

RD-A149 621


Research and Modeling of Supervisory Control Behavior

Report of a Workshop

Thomas B. Sheridan and Robert T. Hennessey, Editors

Committee on Human Factors
Commission on Behavioral and Social Sciences and Education
National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C. 1984



Best Available Copy

This document has been approved
for public release and sale; its
distribution is unlimited.

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

The National Research Council was established by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and of advising the federal government. The Council operates in accordance with general policies determined by the Academy under the authority of its congressional charter of 1863, which establishes the Academy as a private, nonprofit, self-governing membership corporation. The Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in the conduct of their services to the government, the public, and the scientific and engineering communities. It is administered jointly by both Academies and the Institute of Medicine. The National Academy of Engineering and the Institute of Medicine were established in 1964 and 1970, respectively, under the charter of the National Academy of Sciences.

The Committee on Human Factors in the Commission on Behavioral and Social Sciences and Education is sponsored jointly by the Air Force Office of Scientific Research, the Army Research Institute for the Behavioral and Social Sciences, the Office of Naval Research, the National Aeronautics and Space Administration, and the National Science Foundation.

This work relates to Department of the Navy Contract N00014-81-C-0017 issued by the Office of Naval Research under Contract Authority NR 196-167. However, the content does not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred.

The United States government has at least a royalty-free, nonexclusive and irrevocable license throughout the world for government purposes to publish, translate, reproduce, deliver, perform, dispose of, and to authorize others so to do, all or any portion of this work.

Available from: Committee on Human Factors, National Research Council, 2101 Constitution Avenue, Washington, D.C., 20418

WORKSHOP ON RESEARCH AND MODELING
OF SUPERVISORY CONTROL BEHAVIOR

THOMAS B. SHERIDAN (Chair), Department of Mechanical
Engineering, Massachusetts Institute of Technology
RENWICK E. CURRY, NASA Ames Research Center, Moffett
Field, California
BARUCH FISCHHOFF, Decision Research (a branch of
Perceptronics, Inc.), Eugene, Oregon
NEVILLE P. MORAY, Department of Industrial Engineering,
University of Toronto
RICHARD W. PEW, Bolt Beranek & Newman Inc., Cambridge,
Massachusetts
JENS RASMUSSEN, RISO National Laboratory, Roskilde,
Denmark
WILLIAM B. ROUSE, Center for Man-Machine Systems
Research, Georgia Institute of Technology
HENK G. STASSEN, Department of Mechanical Engineering,
Delft University of Technology, Netherlands
DAVID WOODS, Westinghouse Research and Development
Center, Pittsburgh, Pennsylvania

ROBERT T. HENNESSY, Workshop Organizer
JEANNE RICHARDS, Administrative Secretary



COMMITTEE ON HUMAN FACTORS

RICHARD W. PEW (Chair), Bolt Beranek & Newman Inc.,
Cambridge, Massachusetts
NANCY S. ANDERSON, Department of Psychology, University
of Maryland
ALPHONSE CHAPANIS, Department of Psychology, Johns
Hopkins University
BARUCH FISCHHOFF, Decision Research (a branch of
Perceptronic, Inc.), Eugene, Oregon
IRWIN L. GOLDSTEIN, Department of Psychology, University
of Maryland
JULIAN HOCHBERG, Department of Psychology, Columbia
University
K.H. EBERHARD KROEMER, Ergonomics Laboratory, Virginia
Polytechnic Institute and State University
THOMAS K. LANDAUER, Bell Laboratories, Murray Hill, New
Jersey
JUDITH REITMAN OLSON, School of Business Administration,
University of Michigan
THOMAS B. SHERIDAN, Department of Mechanical Engineering,
Massachusetts Institute of Technology
ROBERT C. WILLIGES, Department of Industrial Engineering
and Operations Research, Virginia Polytechnic
Institute and State University

STANLEY DEUTSCH, Study Director
ANNE M. SPRAGUE, Administrative Secretary

CONTENTS

PREFACE	xi
SUMMARY	1
The Concept of Supervisory Control, 1	
The Importance of Models, 2	
Research Issues, 2	
Summary Conclusions, 3	
1 INTRODUCTION	4
2 THE CONCEPTS AND CHARACTERISTICS OF SUPERVISORY CONTROL	8
Definitions, 8	
Levels of Control, 11	
Human Functions, 12	
Goal Setting and Seeking, 14	
Hierarchical Complexity, 14	
Trust, 17	
Analysis of Supervisory Control Systems and Behavior, 19	
Models of Supervisory Control Systems, 20	
Mental Models, 20	
Matching the Mental Models of Operators and Designers, 22	
Assessing System and Mental Models, 22	
Models of Human Performance, 24	
Failure Modes, 26	
3 RESEARCH ON SUPERVISORY CONTROL BEHAVIOR	29
Research Vehicles, 29	

Real Systems, 29	
Simulations, 31	
Laboratory Studies, 34	
Selection of Appropriate Research Vehicles, 35	
Subjects in Supervisory Control Studies, 40	
The Training of Participants in Experiments, 40	
Data Collection During Training, 42	
Individual Styles of Performance, 43	
Implications for Research, 44	
Performance Assessment, 44	
4 THE DESIGN OF SUPERVISORY CONTROL SYSTEMS	48
Special Design Considerations, 48	
The Goal-Directed Character of Supervisory Control, 49	
The Hierarchical Structure of Supervisory Control Systems, 49	
Human Performance, System, and Decision Models, 50	
The Interface Between the Supervisor and the Intelligent Computer System, 50	
Staffing Levels and Policies, 52	
Communication Between Researchers and Designers, 53	
Design Resources, 54	
Accumulation and Sharing Knowledge, 55	
Guidance from Research, 56	
Understanding the Design Process, 57	
5 CONCLUSIONS	59
Characteristics of Supervisory Control, 59	
Analysis and Modeling of Supervisory Control Systems, 60	
Vehicles for Research, 61	
Subjects, Their Styles, and Performance, 62	
Design of Flexible Operator-System Interfaces, 63	
Research Issues Relevant to Design, 64	
Communication Between Researchers and Designers, 65	
REFERENCES	67

FIGURES AND TABLES

- Figure 1 Typical Supervisory Control System, 9
- Figure 2 Five-Step Cycle of Behavior Common to All Supervisory Control Systems Showing Relative Time for Communication and Feedback, 12
- Figure 3 Venn Diagram of a Taxonomy of a Supervisory Control System, 16
- Figure 4 Progressive Complexity of Research Vehicles, 45
-
- Table 1 Matrix to Assess the Relevance of Supervisory Functions to System Control Modes, 13
- Table 2 Considerations in Choosing a Vehicle for Research on Supervisory Control Behavior, 38

PREFACE

The Committee on Human Factors was established in October 1980 by the Commission on Behavioral and Social Sciences and Education of the National Research Council in response to a request by the Office of Naval Research, the Air Force Office of Scientific Research, and the Army Research Institute for the Behavioral and Social Sciences. In addition, its sponsors currently include the National Aeronautics and Space Administration and the National Science Foundation. The committee's principal objectives are to provide new perspectives on theoretical and methodological issues, to identify basic research needed to expand and strengthen the scientific basis of human factors, and to attract scientists both inside and outside the field to perform the needed research. Its overall goal is to provide a solid foundation of research as a base on which effective human factors practices can build.

Human factors issues arise in every domain in which people interact with the products of a technological society. To perform its role effectively, the committee draws on experts from a wide range of scientific and engineering disciplines, including specialists in the fields of psychology, engineering, biomechanics, cognitive sciences, machine intelligence, computer sciences, sociology, and human factors engineering. Experts in additional disciplines also participate in the working groups, workshops, and symposia organized by the committee. Each of these disciplines contributes to the basic data, theory, and methods needed to improve the scientific basis of human factors.

During the past decade the human operator has been assuming a new role relative to technology, namely that of supervisor of an otherwise automated machine, which

in turn may be controlling a vehicle or dynamic process through its own artificial sensors and effectors. Existing models and analytical tools for understanding this new form of human-machine relationship have been found wanting. The added complexity of such a supervisory control system also raises questions about the most effective methods for performing experiments with it, such as what level of simulation is appropriate or whether it is better to collect data from the actual system in operation.

To explore these issues, nine experts met for two days in Sarasota, Florida, in February 1983. The group was charged with answering two questions: (1) How should experiments in supervisory control be carried out so that useful models can be inferred and validated? (2) What does the designer of supervisory control systems need from the researcher, and how can communication between them be effected?

This report provides neither simple answers nor policy recommendations with respect to those questions. The answers we provide are in the context of our discussions on research and design (Chapters 3 and 4), which take up each of the two questions in turn. It is our hope that the efforts initiated by the workshop can be continued by researchers involved in the design and use of supervisory control systems.

In addition to the nine experts, a number of people contributed in important ways to the success of the workshop and to this report. Robert T. Hennessy, the committee's study director in 1983, planned and organized the workshop. Stanley Deutsch, the committee's study director, made valuable contributions in drafting and organizing the report. Christine L. McShane, editor of the Commission on Behavioral and Social Sciences and Education, was extremely helpful in improving the organization, clarity, and style of the report. Elaine McGarraugh provided editorial and production assistance. Jeanne Richards and Anne Sprague provided extensive secretarial and administrative support.

Thomas B. Sheridan, Chair
Workshop on Research and
Modeling in Supervisory
Control Behavior

SUMMARY

Supervisory control is the human activity involved in initiating, monitoring, and adjusting processes in systems that are otherwise automatically controlled. In February 1983, the Committee on Human Factors convened a two-day workshop to recommend approaches to research on this subject and procedures for translating the results into design practice. This report covers three major themes that emerged from the discussions: (1) concepts and characteristics of supervisory control systems, (2) the choice of appropriate research vehicles, and (3) the interchange between researchers and designers.

THE CONCEPT OF SUPERVISORY CONTROL

The term supervisory control generally is used to indicate that one or more human operators are setting initial conditions for, monitoring and intermittently adjusting, and receiving information from a computer that itself closes a control loop (i.e., interconnects) through external sensors, effectors, and the task environment. Another form of supervisory control is involved when a control computer makes complex transformations and integrations of data for display and generates detailed control actions from the operator's commands without immediate feedback to the supervisor. Mediation between the operator and the system processes by an intelligent computer, akin to a knowledgeable human staff member, is the key characteristic of a supervisory control system.

THE IMPORTANCE OF MODELS

Development of explicit models of supervisory control behavior appears to be an effective means of coordination among researchers as well as a means of communication with designers. Models can serve as a common referent for research at all levels, identifying data needs and generating specific predictions that can be tested empirically at all levels, from the real world to the laboratory. Validated models can also serve as the medium for communication with designers, since models provide quantitative behavioral information in a form that designers are accustomed to using.

RESEARCH ISSUES

Researchers can contribute to the design process by first understanding it, then by providing designers with information gained from research in a form that is useful to them. This implies a greater need for communication and coordination between designers and researchers.

There is no shortage of research issues in the field of supervisory control. Research is required at all levels, from the actual systems to the most basic laboratory experiment. At the level of both real systems and complex simulations, research results are likely to take the form of more questions, to be restated as more generalized concepts and examined with greater control in limited or more abstract simulations or in the laboratory. These results in turn must be used in specific applied research and design in order to be validated.

This process of research and validation is expensive and time-consuming but nevertheless necessary, since the principles and theories of supervisory control must be proven to be relevant ultimately in the real environment of automated machine systems. A common view of supervisory control, i.e., a framework for research and for cooperation and coordination of research activities at all levels, is essential. Work at any single level alone will provide neither economical answers to applied problems nor generalized principles that can be used with confidence. Only by spanning the full range of research vehicles, from real systems to the laboratory and back again, will the knowledge gained both be applicable to the design of actual systems and contribute

to a cumulative base of scientific knowledge about supervisory control.

SUMMARY CONCLUSIONS

1. It is useful to characterize the emerging class of human-supervised, computer-controlled systems by strict as well as broader definitions. Salient concepts involve hierarchy, cycles of goal setting and seeking, and trust.

2. No single or simple model of supervisory control is appropriate at this time. Various models have emerged and are useful as paradigms for analysis and experiment. More sophisticated models will be in demand to guide research and design in the future. Supervisory control, while reducing the human operator's participation as a manual controller, depends on human decision-making skills and is vulnerable to error in that decision making.

3. Experimenting with supervisory control systems is difficult for a variety of reasons. Various research vehicles, including real systems, high- and low-fidelity simulations, and laboratory settings, are appropriate at different stages of research.

4. Experienced subjects are essential for research. Researchers must cope with individual operating styles and multiple measures of subject performance.

5. Supervisory control systems can never be completely closed, since the human supervisor must have the capability to set subgoals. The interface of supervisor and computer, especially with regard to high-level cognitive interaction, poses a number of unsolved problems.

6. Better guidance from researchers is needed for designers and operators, in the form of principles and checklists. Better feedback in the form of lessons learned is needed from operators to designers and researchers.

1: INTRODUCTION

Human factors as a professional field is primarily concerned with the compatibility of people and equipment in technological systems. Since World War II, a major emphasis has been on the role of people as system controllers. The design of control systems has gone through dramatic changes in the past 40 years and is now well into the "third wave," to use Toffler's term, of human-machine control systems, (Toffler, 1980).

The first wave was characterized by concentration on the interface itself: the design of displays and controls with emphasis on vision, hearing, anthropometry, etc. The models were simple generalizations of tabular experimental data that were easily adapted to design handbooks (the use and refinement of which is still very much in order). The emphasis of the second wave, in the classical manual control tradition, was on the dynamics of the entire control loop: human and machine are essentially coupled and cannot be analyzed separately. The closed loop block diagram was easily adapted by both researchers and designers. The third wave is brought on, of course, by the advanced technology of the computer and its ability to automatically control, to generate integrated information displays, and to serve as an expert or cognitive aid to operators, on the basis of both prior and current data inputs. The third wave appears to be more extensive and pervasive than the other two.

While we certainly know how to design systems incorporating the technology of the first two waves, we are neophytes in knowing how to design systems in which a person nominally directs and oversees processes

controlled by a computer, an interaction known as supervisory control. We know little about what strategies are or might be used by the human operator or the computer, how to relate available resources to control demands, how best to allocate functions to people and computers, or how errors arise in the interaction. Nevertheless, the very availability of the new computer-based technology seems to have set its own imperative that both government and industry adopt it, whether or not the art and science of design are ready.

Supervisory control occurs in a variety of systems, for example, conventional or nuclear power plants, propulsion systems, modern aircraft, and command, control, and communication systems. Supervisory control behavior has been the subject of numerous research projects, and there has been considerable interest and effort in the direction of modeling this category of control behavior, using extensions of human performance modeling. However, since most of this work has focused on a particular system, there is a need to broaden the conceptual understanding and to develop general principles for the practical aspects of job design for human controllers. In recognition of this need, the committee devoted one of the six chapters of its report Research Needs for Human Factors (National Research Council, 1983) to discussing a broad spectrum of research to provide the data and principles for design of future supervisory control systems.

There is serious controversy over the research strategies appropriate for investigating supervisory control issues and how to approach modeling of supervisory control behavior. And there is no established protocol for verifying and validating research and modeling results through comparison with supervisory control performance in actual systems.

There appear to be three alternative approaches to the study of supervisory control behavior: (1) to study actual systems operations, e.g., power plant operations or command and control centers, by monitoring and analyzing everyday, complicated tasks and events and applying principles and lessons learned to new situations, (2) to develop and use large-scale simulations, with all the attendant problems of cost, personnel training, and operation, and (3) to adopt or develop simplified paradigms suitable for small-scale laboratory studies, with the risk of missing critical

conditions and the difficulty of generalizing results to full-scale operations.

The workshop that is the basis of this report continued the theme of the chapter in the committee's report; we attempted to identify the research that would improve our understanding of supervisory control behavior. The principal topics we discussed were how the research could be accomplished and how the results could be translated into design practice. Stated in the form of questions, the workshop had two objectives:

1. How should experiments in supervisory control be carried out so that useful models can be inferred and validated? Should real (usually complex) systems be observed directly, even though experimental control is not practicable? Should real systems be simulated in relatively high fidelity, to permit controlled experiments in which malfunctions, such as overloads and failures, are forced to happen? Should experiments be downscaled to much simpler tasks that are somehow still analogous? Should all these approaches be pursued? How can investigators in the area share facilities and data?

2. What can researchers offer to designers of supervisory control systems, and how should the transfer occur? Given that cognitive and supervisory control research results and models are (and are likely to become) more complex and more qualitative, how do we move from this research/modeling domain to the design domain? That is, what design recommendations can we make to best serve designers' needs?

Each participant prepared a brief paper on these questions, which served as starting points for discussion at the workshop. These papers, synthesized with the results of our discussions, form the contents of this report. We attempted to capture all the ideas expressed and organize them in a coherent manner. Several topics are treated in more than one place, (e.g., goal setting and seeking), and some topics are given greater emphasis than others, (e.g., subjects in experiments involving performance measurement). Our intent is to reflect the character of the discussions at the workshop rather than to provide a tightly organized and balanced treatment of all issues in conducting research on supervisory control behavior.

The report is divided into five chapters: Chapter 1 is an introduction to the report. Chapter 2 presents

the concepts and characteristics of supervisory control systems. Chapter 3 addresses research needs to better understand and predict the behavior of people functioning as system supervisors, including the issue of which research vehicles (i.e., real-world, simulation, or laboratory) are appropriate for different research purposes and the related issues of choosing subjects and measuring performance. Chapter 4 discusses the development of fundamental principles of design of supervisory control systems; it deals specifically with communication between researchers and designers, areas in which researchers can be helpful to designers, and some thoughts on the nature of the design process. This chapter explicitly recognizes that supervisory control behavior, whatever its intrinsic interest as a topic for basic research, has become a matter of concern because of its implications for the design and function of actual systems. Chapter 5 presents a number of conclusions that emerged from the workshop discussions.

It is important for researchers to recognize that the interest in understanding supervisory control behavior is to effect practical improvements in the design of future systems. If they are to contribute to these improvements, they must be aware of the problems and constraints faced by designers, the kind of information they need, and the nature of the design process.

2: THE CONCEPTS AND CHARACTERISTICS OF SUPERVISORY CONTROL

DEFINITIONS

Simply stated, supervisory control refers to all the activities of the human supervisor who interacts via a computer with a complex and semiautomatic process. The term supervisory control is derived from the close analogy between the characteristics of a supervisor's interaction with subordinate human staff members and interaction with automated subsystems. A supervisor of people gives general directives that are understood and translated into detailed actions by staff members. In turn, staff members aggregate and transform detailed information about process results into summary form for the supervisor. The degree of intelligence of staff members determines the level of involvement of their supervisor in the process. Automated subsystems permit the same sort of interaction to occur between a human supervisor and the process (Ferrell and Sheridan, 1967). As the committee's report discusses (National Research Council, 1983), supervisory control behavior is interpreted to apply broadly to vehicle control (aircraft and spacecraft, ships, undersea vehicles), continuous process control (oil, chemicals, power generation), and robots and discrete task machines (manufacturing, space, undersea mining).

In the strictest sense, the term supervisory control indicates that one or more human operators are setting initial conditions for, intermittently adjusting, and receiving information from a computer that closes a control loop (i.e. interconnects) through external sensors, effectors, and the task environment (Figure 1).

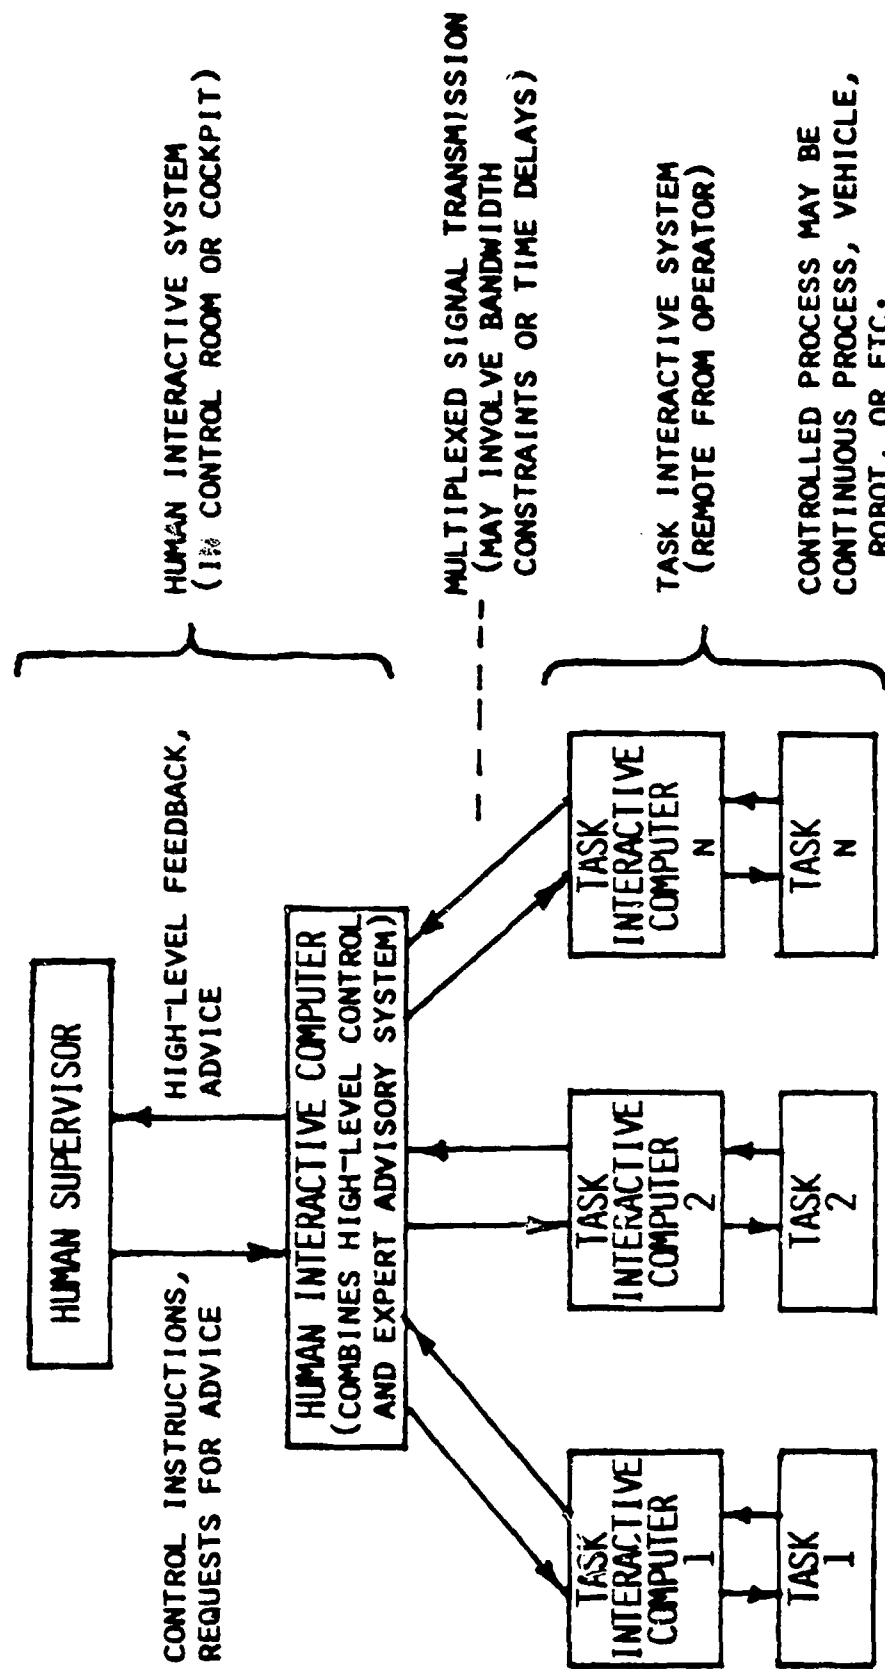


FIGURE 1 Typical Supervisory Control System

In a broader sense, supervisory control is involved when a computer makes a complex transformation of data to produce integrated (chunked) displays, or retransforms operator commands to generate detailed control actions even without immediate feedback.

The essential difference between these two characterizations of supervisory control is that in the first and stricter definition the computer can act on new information relatively independently of and with only blanket authorization and adjustment from the supervisor; that is, the computer implements discrete sets of instructions by closing the loop through the environment (completing a causal chain). In the second definition the computer's detailed implementation is open loop. The two situations may appear similar to the supervisor, since he or she always sees and acts through the computer (analogous to a staff) and therefore may not know whether it is acting open-loop or closed-loop in its fine behavior. In either case the computer may function principally to implement the supervisor's commands, principally to interpret incoming information from below and give advice to the supervisor, or both.

Two of the principal characteristics of a supervisory control system are semiautonomous action and complexity. Consequently, the human supervisor cannot simultaneously be aware of all events occurring in the system. Moreover, the supervisor may not have the capacity to assimilate and assess all factors relevant to making a control decision, even if he or she had access to all the necessary information. Therefore, to permit a person to act as a competent supervisor, the system ideally would be designed to (1) provide all the information than is appropriate for a particular decision, but no more, (2) provide it in the most understandable form, (3) alert the supervisor to conditions that may require attention, (e.g., failures), and (4) aid the supervisor by suggesting possible courses of action or at least laying out the likely results of the alternatives.

These ideal characteristics require intelligence in the system. For a system to act intelligently, both to aid the supervisor and control its own semiautonomous functions, it must have an internal model of its own structure, logic, and dynamics. In order to support strategies chosen by a supervising operator, an intelligent, assisting computer must also have a model of the operator's decision style and preferences. For example, by maintaining data of its operating history,

the system can provide part of a knowledge base to aid operators in various ways, such as helping them learn from their own or other operators' experience or ask questions about past history, i.e., what happened when. It can also help evaluate future possibilities, i.e., what would happen if ____ (which implies running the computer model as a dynamic simulation, possibly in fast-time, with initial conditions and control or disturbance inputs specified by the operator). This same model and data base would be used by the computer for automatic control and for diagnosing failure.

Levels of Control

It is important to note that currently there are at least two levels of control, each involving a different computer with different functions: the supervisory level, called the human interactive subsystem (HIS), and the subordinate level, called the semiautonomous task interactive subsystem (TIS), where specific tasks are controlled (Sheridan, 1984a). The HIS computer may be more or less sophisticated in understanding and implementing commands, assessing situations, or giving advice. One supervisor-plus-computer (HIS) may serve many low-level control systems (TIS). It is especially important to emphasize this fanning out or multiplexing of control at the lower levels, a proven principle of organization for both physical and biological systems. Figure 1 illustrates the structure of levels of control, the couplings and the fanning out at the TIS and task levels.

In the simplest case supervisory control may involve the simple human decision and action to override or modify one automatic mode for another; for example, an elevator operator, observing that a passenger is about to enter the elevator, inhibits and reverses the automatic closing by pushing the "door open" button. Here there is little or no sophistication required in the communication between human supervisor and computer (in this case relay logic). In more sophisticated forms of supervisory control the human interactive computer may have elaborate means for advising the operator (involving text or graphics or both) or for understanding the operator's queries or commands. Indeed, the branching or hierarchical structure itself may be recursive, with more levels of computational

supervision than is shown in Figure 1. In any case, the behavior of the human supervisor tends to be intermittent more than continuous, cognitive more than perceptual-motor.

Human Functions

In all forms of supervisory control there is a typical five-step cycle in the human supervisor's behavior (Figure 2): (1) planning, including the setting of subgoals relative to the given task goals, (2) instructing the computer, (3) monitoring its execution of instructions and making minor adjustments, (4) intervening to circumvent the automatic controller as necessary, and (5) learning from the experience in order to plan better (Sheridan, 1984a). Iterative feedback and communication usually occurs between learning and planning at long intervals, between intervening and instructing the computer at intermediate levels, and at very short intervals within the monitoring and adjusting step.

The importance of these functions for the process being controlled must be evaluated in the context of the particular system (see Table 1). Every process has four possible control modes: (1) normal start/stop, (2) normal process operation, including automatic control

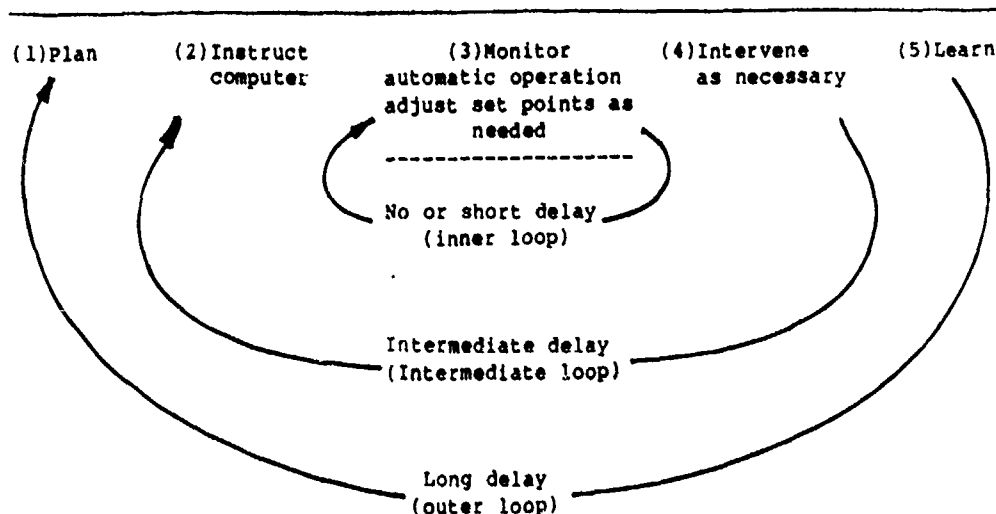


FIGURE 2 Five-Step Cycle of Human Supervisory Behavior Showing Relative Time for Communication and Feedback

TABLE 1 Matrix to Assess the Relevance of Supervisory Functions to System Control Modes

Human Supervisory Functions	System Control Mode			
	Normal Start/Stop	Normal Process Operation	Fault Management During Operation	Maintenance (nonoperation)
Planning	*	*	-	-
Instructing	*	*	-	-
Monitoring	*	*	*	*
Intervening	*	-	*	*
Learning	*	*	*	-

Planning (primarily for start/stop) is deciding what to instruct the computer to control automatically and when to shut down.

Instructing is programming plans into the computer to do (or start to do) certain things automatically for normal operation, or to stop some actions when they are completed or abnormal.

Monitoring is watching the (usually) normal automatic operation of the system to see if it is satisfactory and manually adjusting set points of automatic controllers as automatic control continues.

Intervening is breaking into the automatic control loop to stop one task and start a new one, to take emergency actions (fault management) manually, or for maintenance or repair.

Learning is gleaning from experience what is necessary for better planning or other supervisory functions.

Note with reference to Figure 2 that normal system operation, including set point adjustment, occurs within the inner loop; minor normal start/stop or fault management occurs within the intermediate loop; and major start/stop and fault management occurs within the outer loop.

and process tuning, (3) fault or abnormality management during operation, and (4) nonoperation for maintenance, repair, or lack of demand. While all process systems have supervisory functions, there are marked differences among them as to whether a particular function is performed primarily by the human supervisor or by the computer. For example, in nuclear power plants, the primary fault management or "engineered safeguards system" (boron rod insertion, high-pressure coolant injection) is triggered electronically for extreme abnormalities, but the human supervisor must take follow-up actions and, in the case of less threatening failures, manage them entirely. For most complex technological processes, it is currently believed that the human supervisor is far better at optimizing in nonroutine circumstances than any computer-based optimal control system. This is in part because we are unable to quantify the criteria for control and equations for the system, especially in nonroutine situations.

Goal Setting and Seeking

One way to describe supervisory functions is to analyze the system in terms of the interplay between goals and means: given a goal, what means or mechanisms support that goal; what goals are affected by a given means. The result is a mapping of the problem space within which supervisory activities occur. Goals and specifications propagate from the top down; resources and limitations come from the bottom up. Control and decision tasks (what is to be done) can be formulated for a process at each level; the reasons for decisions are found at the next higher level; and the resources for implementing a process are sought at the next lower level (Rasmussen and Lind, 1981). It would be useful to develop guidelines to identify the kinds of decisions that are appropriate at each level of the control task.

Hierarchical Complexity

Two subsystems in a supervisory control system are at different levels in the hierarchy if control passes unidirectionally between them. Several observations can be made about the nature of the complexity of systems organized in this manner. Complexity is a function both

of the number of levels and the number of subsystems at any level. Complexity is also an increasing function of the time taken for information or control to pass from its source to its destination and the uncertainty (in an information theoretic sense) in the relation between subsystems. Finally, complexity is an increasing function of the amount of coupling between or among subsystems.

Several conjectures can be made about how complexity influences the functions of the supervisor and his or her perceptions of the system's characteristics: first, the greater the difference in levels through which control is organized, the wider the span of control, the longer the time constant, the more wide-ranging the effect of control, the greater the complexity experienced, the greater the probability of loss of control due to inappropriate action--unless subordinates exercise local control by interpreting global commands. Second, perceived complexity (and hence load, difficulty, etc.) is a direct function of the number of transformations (rule to action, signal to symbol, etc.) that must be performed by the operator. It is desirable that perceived complexity be low.

Another factor that tends to make supervisory control systems complex is the nature of their hierarchical structures. In such systems, relative to any particular level in the hierarchy, activities at other levels tend to be loosely coupled and consequently obscured. During normal operations, this separation can be efficient, but during abnormal operations this can make it very difficult for the supervisor to assess the state or configuration of the system. Appropriate aids must be designed to deal specifically with the complexity of abnormal operations.

Figure 3 is a Venn diagram describing the interactions between the human supervisor, the computer, the semiautomatic controllers, and the slave controllers. This figure suggests problems that may be expected to arise if the human supervisor only, the HIS computer only, or neither is sufficiently connected to the task itself. These problems appear as phrases connected with arrows to the task interaction areas, which have different degrees of connection to the HIS computer and the human supervisor.

The ability to predict the consequences of actions initiated by either a human supervisor or an automatic controller is fundamental to the notion of control. In




FIGURE 3 Venn Diagram of a Taxonomy of a Supervisory Control System

Source: Boff et al. (1984).

a supervisory control system, this predictive ability is likely to reside in both a computer-based model and in the human operator's mental model. There are regions of interaction in which one influences the other. There are other regions in which no interaction is possible

between HIS and TIS or in which possible interactions may degrade system performance.

It is unlikely that the complexity of a supervisory control system could be defined independently of both system context and of the contextual knowledge of the people and computers involved in the system. This issue needs attention from design-oriented researchers in human-machine systems.

Trust

Supervisory control demands that the system be trustworthy. What the human supervisor can find out about the state of a system, particularly during abnormal or emergency situations, will greatly influence whether he or she trusts that the lower-level control system will carry out a command or meet an objective that is passed down. If such trust is not there, even if warranted but not perceived, the operator may spend much additional time checking low-level system functions. Note that it is the perception of trustworthiness that determines operator behavior, not actual trustworthiness. Beyond some point, the operator will prefer to do the task himself or herself--and indeed be more efficient at it--without the "help" of the computer and automation.

The question of trust has two dimensions. One is the predictability of the consequences of different actions. The second is their desirability. If trouble cannot be avoided, one would at least like to be able to predict it. For operators to make the correct allocation of functions between themselves and the system, they need an accurate perception of its trustworthiness. In practice, that perception expresses itself in judgments regarding both the system's performance and its reports on its own performance, such as: Can I trust those warning indicators? Will this emergency subsystem work if put to the test? Can I push it a bit further? What will happen if maintenance is delayed? A measure of an operator's trust in the system (e.g., whether he or she diagnoses situations as emergencies) would facilitate predicting that person's performance. If there is some systematic bias in that trust, then the operator might benefit from better system diagnostics or training. The appropriateness of trust (like the appropriateness of other aspects of an operator's perception of the system)

should increase with time, assuming that the system provides some reasonable feedback. If that is not the case, some aspects of its performance need to be made more transparent.

The expected behavior of the novice operator is not to trust a new system or ill-understood situation--which is usually appropriate, at least in part, because novices are themselves unreliable. Greater transparency to the supervisor of the workings or performance of a system will not necessarily result in greater trust. However, there is conflicting evidence on this point. Wood (1982) has shown that relatively novice operators, when offered the services of a computer algorithm to perform a given task instead of doing the task themselves, may choose the computer even though readily available evidence shows that they can do it better and faster themselves.

Transparency implies that users can "see through" to the workings of the system, that their perception of the task situation is correct, despite displays presenting information that does not match the concepts they normally use.

This is not to say that some degree of integration of information at the level of displays and controls is not a good idea or is not dictated by other requirements. For example, when operators simply cannot cope with too many details, they need some summarizing or abstracting at an appropriate level of abstraction. The best solution is probably the provision of alternate display-control modes--some that allow more direct communication with details of the task and others that amount to communication through an intelligent intermediary, an honest and trustworthy broker.

The appropriate criteria for trust need to be studied to develop a theory of trust in supervisory control. Trust may have many ingredients; at least in part it has to do with statistical confidence in what a system element will do and the expected consequences. In the same vein might be considered the trust that the designer places in the computer and the system as well as in the operator. The displays, controls, and operating procedures are designed according to the designer's degree of trust or interest in these elements.

ANALYSIS OF SUPERVISORY CONTROL SYSTEMS AND BEHAVIOR

Rasmussen and Lind (1981) use the terms purpose, function, and equipment to distinguish the levels of analysis of a complex system. Because supervisory control systems are multilevel systems, it is very difficult to develop a coherent structure of the levels without such analysis. Although in rare, unfamiliar situations, control actions cannot be preplanned in detail, descriptions of disturbed functional states can be formulated in terms of discrepancies from what is acceptable (permissible target states) at the various levels of the purpose-function-equipment hierarchy. These discrepancies then define the control requirements. Top-down analysis is necessary for formulation of target states in the hierarchy, and bottom-up analysis is necessary to identify the relevant disturbed states and the range of their possible values.

Designers of supervisory control systems might aspire to meet the design requirements of formal control theory (knowing that such objectivity and quantification are not possible in practice). Formal control design requires that the task or controlled process be specified in terms of equations relating outputs to inputs and expected disturbances. In addition, the "objective function" must be stated quantitatively, i.e., how the goodness (quality) of performance is explicitly related to error (system states achieved relative to those desired) as well as time, energy, or other economic resources consumed. Finally, the means available, both hardware and personnel, to control the process must be specified. Formal control theory then specifies how to optimize the objectives (maximize goodness) in terms of control laws, or come as close as possible with given resources.

Human supervisory control will remain deficient against this theoretical yardstick for several reasons: (1) the real controlled process is never fully understood, (2) there are usually as many "objective functions" as there are people, none of them specifiable in other than "fuzzy" (linguistic) terms, and (3) human implementation is severely constrained in speed, power, and precision.

MODELS OF SUPERVISORY CONTROL SYSTEMS

Although a variety of models of supervisory control have been proposed (Sheridan and Johanness, 1976; Baron et al., 1982), there is little consensus as yet. A model can be a paper description of a system; it can be a functional model implemented on a computer, which emulates the function of a system; it can also be a mental model, an internal representation of a system held in the mind of an operator, designer, or researcher. In all cases the intent is to represent the functional and physical features of the system that are salient to the decision task at hand. These models differ in kind and detail.

Learning to operate a system usually includes not only learning fixed procedures, but also learning something about the nature of the system, so that inferences can be drawn and actions taken on the basis of what the operator "knows" about the system. Even after considerable experience and effort in both of these types of learning, it is unlikely that he or she will have been explicitly taught what actions are appropriate under all possible circumstances.

Just as an operator's mental model may differ significantly from the actual physical and functional characteristics of the system, similarly a computer-based model of a system may also differ significantly from both the system itself and the mental model of the operator. (Even the designer's model of a system may not completely match the system; properties and functional interrelationships of the physical system may lead to operational consequences in both normal and abnormal situations unanticipated by the designer.) A central issue in research and design of supervisory control systems is to bring all these models into harmony, since they ultimately influence the decision processes of the human supervisor and the consequences of the system's operation.

Mental Models

A major research problem is how to measure or otherwise infer what an operator's mental model is. There are formal subjective assessment techniques, such as interpretive structural modeling (Warfield, 1973), policy capturing (Hammond et al., 1964; Slovic and

Lichtenstein, 1971; Dawes, 1979), multiattribute utility (Keeney and Raiffa, 1976), and fuzzy set theory (Zadeh, 1965; Sheridan, 1984b). Most such formal techniques are designed to elicit information in a very specific form--which may not measure the relevant information in the mental model or the aspects of it that might eventually be at odds with how a system was designed. Less formal techniques include debriefing, recollection of critical incidents, verbal protocols (Ericsson and Simon, 1984), and qualitative modeling (Gentner and Stevens, 1983). All of the latter are more or less similar, being free-form discussions in which a knowledgeable interviewer poses questions or goes through a checklist to elicit judgments about task structure, priorities, probabilities, consequences, etc. Sometimes replaying videotapes or showing objective data can help the subject remember and reflect on key situations.

The problem of revealing an operator's mental model of a system is exacerbated by the fact that often an individual may be unable to describe this knowledge even to himself or herself. Moreover, an operator's description may be contradicted by what he or she actually does.

One means of investigating mental models other than direct questioning about them is to assess the operators' perceptions of system contingencies (e.g., what will happen if I do ___). This approach offers one way to get a general measure of the accuracy of their perceptions--of how well they understand the limits of their own knowledge. An ambitious way to do that would be to request of them explicit probabilistic predictions of the likelihood of each possible response of the system to each possible action. Once the outcome is known (which may be immediately or only after careful reconstruction), we can see how well operators understand the system and how well they appraise their own understanding. If many confident predictions prove to be in error, then an operator is overconfident and is in need of better training or information. These evaluations can also be fed back to operators. For some 15 years the U.S. Weather Service has collected such probabilistic judgments and provided forecasters with modest performance feedback. The objectives is to improve forecasts, provide more precise information to users of those forecasts, and identify information needs (Fischhoff, 1982).

Matching the Mental Models of Operators and Designers

Matching the mental model of operators to models of the system could be thought of as matching the mental model of operators to the mental model of designers. From that perspective, a crucial problem for research is to discover what designers know about how operators think and act. This problem could be studied by showing designers videotapes of actual operations that are stopped periodically with a request to predict what happened next, or by having designers predict operators' learning curves. To improve communication between operators and designers, by whatever institutional measures are available, should be a major goal. Designers could be required to provide a model of human response with their models of the physical system. It has been proposed (Goodstein, 1982) that operators be provided with in-situ information as to why designers designed certain equipment or procedures the way they did, in addition to information about what the equipment is and how it works. Such information can be crucial to operators in judging whether it is acceptable to modify procedures or interfere with interlocks during emergency situations. At present, such situations often require that designers be consulted.

Alternatively, improving the match between the mental model of the operator and the system model might increase the speed with which operators learn how the system operates. From that perspective, an important problem for research is to discover the operator's image of the system and then to assess how its accuracy changes over time. The key design tasks are to make the system's operation more transparent and provide better feedback to operators.

Assessing System and Mental Models

As with all others, decisions in a supervisory control system involve choices among possible courses of action--e.g., continue operation, take emergency action, adjust the set point of an automatic controller, do nothing, collect information, or consult with others. If it is possible to identify the components of decision problems according to the analyses discussed above, then it is also possible to identify the best course of action and thus to evaluate the operator's performance. The key

components are: the set of possible actions, beliefs about the state of the system and the effect of different actions, and the set of goals (with some idea of their relative importance). Should an action seem inappropriate, then it may be due to neglect of relevant action possibilities, misunderstanding of how the system is operating (and how it would be affected by various interventions), or misinterpretation of what goals are pertinent, for example, because information on design rationales is not available. Each of those problems calls for a somewhat different solution; for example, better diagnostic displays can reduce confusion about what is happening, but not confusion about what should be done when there are conflicting goals.

It is not always possible to elicit the full decision problem that an operator is considering. Time may be too short to ask many questions, actions may not be chosen in a deliberative fashion, and it may be hard for the operator to articulate the reasons for action. However, consideration of even simplified and implicit decision models can be useful in a number of contexts.

One such context is design review. At each stage of a system's operation it should be possible for the operator to identify the best course of action. If that cannot be done, even in principle, then some redesign is called for. Consider several examples. If the time is too limited to access all critical information, then a better display is needed. If some of the systems' signals are likely to be misleading or if disagreements within the operating team are possible, then special training may help. If goals conflict (e.g., should I do this operation manually in order to preserve my skills or should I allow it to be done automatically for the sake of efficiency), the organization responsible for the system needs to develop an explicit policy (and adhere to it, even when evaluating operators).

In situations that allow for it, comparing operators' perceptions of the decision problem with what the system's designers (or the operators' supervisors) imagine those perceptions to be may be useful. Doing so may reveal when actions that seem obvious to designers never occur to operators, or when operators are wrestling with conflicts that never occur to designers. In effect this procedure contrasts operators' and designers' (or supervisors') mental models of the system in situations that might require action. The models are stylized in the sense of capturing perceptions at a particular point

in time and in reducing rich mental models of system operation to summary predictions of what will happen if ___ (I do nothing, I shut it down, I call for help). They do allow a wide variety of situations to be described in a common set of terms for actions, expectations, and goals that is close to natural language so that people should learn it quickly and use it easily.

This section has concentrated on issues concerning system models, mental models, the decision processes of the human supervisor, and their interrelationships. In a few succinct statements we can summarize why these are important concepts and what they imply for the design of supervisory control systems.

- o If an appropriate mental model is acquired by the operator, then even if the system is complex the operator will be able to handle it. Indeed, only if a mental model is efficiently incorporated to the point at which the operator does not have to think consciously about what is going on will the operator be able to handle very complex relations.
- o To the degree that perceived information matches the model, the operator's behavior will tend to be automatic (skilled, skill-based) and efficient.
- o Training and design should aim to instill the appropriate model in the operator.
- o Displays should service the mental model.

Models of Human Performance

Models of human performance in supervisory control behavior may be one means for both guiding research needs and communicating the results in a form useful to designers. Reviewing the present state of the art of the supervisory models reported in the literature, we can draw three conclusions:

- o For the three process control modes of start/stop procedures, process tuning, and fault management, only separated qualitative models of human behavior have been developed.
- o All reported models describe special human supervisory functions such as monitoring, learning, planning, decision making, set point control, or intervention.

- o None of the available models has been applied to more than one of the process control modes, and none of the models is able to describe human supervisory control behavior in an integrated way in terms of the special functions just mentioned.

In terms of readily applicable and predictive models in concise form, the results of a decade of supervisory behavior modeling are not very promising. This does not imply that work on modeling human supervisory behavior is pointless and that research on this topic is useless--just the opposite is true. We are learning to see supervisory control not as an elemental class of behavior but as inherently integrative. Thus much more is demanded of a model of supervisory control than of a model of manual control. There is hope that future research will provide answers to the complex questions involved in modeling supervisory control behavior.

Rasmussen's three-level model of human behavior (Rasmussen, 1983) is useful in thinking about modeling supervisory behavior. In brief, the lowest level is skill-based behavior, akin to simple, servo-mechanistic control. The intermediate level is rule-based behavior, characterized by particular stimulus conditions eliciting particular sequences of actions. The highest level is knowledge-based behavior, at which actions result from assessing circumstances with respect to goals and weighing alternative strategies. This paradigm is strong and useful, although it is qualitative in nature. In general we can say that only skill-based and rule-based behavior can now be modeled, whereas knowledge-based behavior currently evades modeling.

Despite the lack of knowledge in this domain, models of possible problem-solving strategies, together with psychological models of criteria for choice, can be very effective for design purposes. In modeling supervisory behavior, well-defined tasks can be allocated to the computer and ill-defined tasks to the human operator. This important use of models seems more desirable than the current practice of process and control engineers of trying to automate whatever possible.

FAILURE MODES

Systems, including supervisory control systems, can fail in many ways, from physical breakdown to human error. Failure can be immediate, gradual, catastrophic, or simply lead to inefficiencies in the system process. Some failures can be anticipated, with backup or recovery methods preplanned, and others cannot. Failure can occur due to a single event or a sequence of events. All these considerations are important to understanding system failures. Of concern to us, however, are failures that are unique (or nearly so) to supervisory control systems, in particular, those that are attributable to the human supervisor or that might be avoided through his or her decisions and actions.

One of the principal roles of people in supervisory control systems is to serve as a backup when automatic systems fail or are unable to cope with unanticipated situations. While people are expected to perform well during normal operations, the implied responsibility to not make errors is especially great when they are the final authority or last resort for taking appropriate action to avoid or recover from system failure. Under these circumstances, with their many attendant uncertainties, there is the greatest potential for decision errors.

Human failures can be classified into several causal categories:

- o Adaptation is required beyond the capability of the operator, i.e., in terms of time, knowledge, processing capacity, or precision of execution;
- o Slips, in which the operator's intention was not fulfilled; and
- o Mistakes, in which the operator's intention is later considered inappropriate. Mistakes break down into several categories:
 - Functional fixations, interference from highly trained or familiar patterns,
 - Mistakes due to conflicting, overlapping, or similar relationships in the purpose-function-equipment hierarchy, i.e., designed-in traps, and
 - Inappropriate actions to get unneeded information or to test hypotheses unnecessarily.

These types of failure are associated with somewhat immediate events and processes of information assimilation, decision, and control. In general, it is more effective to avoid these failures by adequate system analysis and design, rather than to attempt to eliminate them by better training or by trying harder. However, the goal should not be to completely preclude some modicum of variability of behavior needed for learning. High skill and know-how are developed by trying to do what one has not done before, by cutting corners and by finding short-cuts. This leads to errors only if the system is totally rigid, i.e., an unfriendly system. Human error results from both too infrequent or too frequent control.

A second sort of human failure has origins more remote than the immediate circumstances of controlling the process. These failures are associated with the human operator's knowledge of the system, past experiences, training, and style of operation. Such operator-induced failures include:

- o Control exercised too infrequently or lacking force relative to the system's divergence rate or convergence rate to desired states and
- o Control exercised too frequently or too forcefully, with consequent noise injected into the system and resulting instability.

The operator may exercise control too infrequently because the system appears trustworthy (i.e., it has always done well in the past), because the operator cannot cope with the situation and may wish to avoid responsibility, or because he or she forgets. The operator may exercise control too frequently because he or she does not trust the system (i.e., it has failed in the past) or does not understand what the system will do if left alone, or because the operator likes to be involved. The correct magnitude of control may simply go unknown because of delay in feedback.

Preparing and assisting the operator to cope with unexpected or emergency events is one of the central problems of research on supervisory control systems. It involves training procedures and work station design. The handling of most routine events can be automated or done by a person with any of a variety of control interfaces. Yet as is well known, emergency events, even near-emergency events, are both rare and unique.

The more complex the system, the more difficult it is to arrive at meaningful categories into which events and appropriate control decisions may be classed. In a real emergency, the operator or crew takes only one of many possible routes through the tortuous path of decision logic that unfolds during the critical event. The appropriateness of alternative paths and their implications will always have some uncertainty.

Most human-involved failures in supervisory control systems can be viewed as decision errors, ranging from procedural slips to choosing an inappropriate level of involvement in control of the system. While decision errors continue to be a problem in many types of systems, a new class of decision problems may arise in particular with regard to supervisory control systems. These problems will depend on the characteristics of the interfaces within the total decision system and their consequent effects on the elements of trust, skill maintenance, and information accessibility.

3: RESEARCH ON SUPERVISORY CONTROL BEHAVIOR

RESEARCH VEHICLES

In general, research on supervisory control behavior serves two purposes: to develop behavioral principles to serve as a basis for design and to evaluate specific design solutions. Empirical data on supervisory control behavior can be obtained from analysis of actual incidents in real systems or from simulated incidents. Simulated tasks may be of high fidelity (attempting to recreate the work environment of interest), low fidelity (abstracting the critical elements into a generalized work environment), or laboratory studies. We discuss below research supported by a number of vehicles: (1) real systems; (2) high-fidelity, comprehensive simulations; (3) low-fidelity, limited simulations, and (4) context-free laboratory studies. The appropriate vehicle depends on the nature of the research. Each data source has advantages and limitations; probably a multiple method or "converging operations approach" is ideal for progress in this area.

Real Systems

In general, real systems are both beginning and end points for research on supervisory control behavior. They determine the scope and the nature of the problem set, provide information on task requirements and actual operator behavior, and are the final proving ground for the design, training, and procedural concepts originating and developed in laboratory and simulator settings. For

example, operational systems are the best vehicle for analyzing behavior to identify such things as the subjective goals of operators, the information-processing strategies actually used as a function of experience and training, criteria controlling the choice of strategies, e.g., time, trust, and cognitive strain, possible strategies, habitual short-cuts, work and team organization and constraints on changes in organization, the influence of company policies, and the effects of typical disturbances, e.g., interference from management or government authorities during critical events.

A limitation of research using real systems as the vehicle in process control is that no two physical plants are alike. Perhaps more important, the plants tend to differ in operating philosophy, management style, and training procedures, not to mention in the extent of cooperation between engineering and operating personnel. Extremely careful analysis and good judgment are necessary to pick a level for which apparent comparability across circumstances is valid. (These individual differences in systems may not exist to the same extent for aircraft or other systems).

However, in complex systems generally, it is very time-consuming and costly to collect the data needed to analyze and understand the operator's decision logic in a particular event. It requires documented evidence of what happened; it requires interviews with the operators promptly following the critical events, in which they must relate their recollections to the factual record of events. Standard reports on critical incidents do not routinely contain the data necessary to characterize human-machine system performance. Retrospective verbal reports from those involved in these incidents can provide insight into supervisory control behavior provided techniques to analyze the decision process are used (e.g., Pew et al., 1981). Even so, the task can be arduous.

For example, with the reasonably complete, traditional accident analysis reports available, it took nearly 120 hours of work to capture and analyze the decision-making details of each event in a study of failures in nuclear plants (Pew et al., 1981). The investigators had access to only four major events; a far larger number would be needed to ensure that the general conclusions drawn would be valid. Further work is needed to refine these procedures and to make them a standard part of incident analyses. Because retro-

spective verbal reports are subject to a variety of biases, successful analysis of this type of data depends on corroborating data on actual system performance and operator action during an incident. However, in many industries performance data are fragmentary; data logging systems are needed if analysis of actual incidents is to provide data on human performance in supervisory systems.

Data collection should be built into complex supervisory control systems so that operator errors as well as operator error recoveries can be recorded (they should get credit for the recoveries). Since attribution of blame is not important for research, researchers should aggregate errors so that individual attribution of error is difficult. Perhaps even more important is to record the types of messages (encoding) used, cognitive aids and intelligent capabilities employed, and timing, especially in emergencies. These are important to observe since real supervisory control systems are so flexible and unpredictable and experimental control is simply impractical. This kind of analysis could become practical only if a very well-structured procedure could be developed that greatly streamlined the process--and that requires a theory of supervisory control.

Simulations

Simulators as research vehicles have several advantages: conditions and events, particularly rare events such as emergencies, can be structured and controlled, scenarios can be repeated, detailed data can be collected, and usually the cost of research is much lower than for investigations involving real systems. Of course, when rare events are duplicated in simulators, they are no longer rare events, and the behavior of subjects may be affected.

For any simulation study, especially when high-fidelity simulation is used, much thought should be given to the characteristics of subjects, (i.e., the transfer of training, positive or negative, from their normal work to the experimental tasks and the amount of pretraining required) as well as the need for constructing realistic scenarios, including management philosophy, etc. Careful consideration should also be given to requirements of data logging and analysis and

to ancillary questions, e.g., would it be useful to collect protocols and, if so, will this affect performance. A frequently overlooked point in research planning is the possibility of drawing on other ongoing work to determine whether, at relatively slight cost, additional questions could be asked that would increase the generality of the findings.

High-Fidelity Simulation

High-fidelity simulations have the strengths of affording close control of the test situation, the opportunity to simulate rarely occurring faults and conditions in a realistic world configuration, and the feasibility of very systematic and complete data collection, measurement, and processing. They are excellent vehicles with which to identify problems in information retrieval and search functions in complex display systems, to evaluate capacity requirements during real-life scenarios with complete crews of operators, and to evaluate the match between the operators' choices of available displays and strategies and the designer's intentions. High-fidelity simulation is also an excellent means for teaching skill-based and rule-based behavior. (Whether it is also justified for teaching knowledge-based behavior is uncertain, despite the fact that this is often purported to be the objective of high-fidelity training simulators. Simple "concepts trainers" seem to be preferable for initial training.)

There are certain disadvantages to high-fidelity simulation as a research vehicle. While controlled scenarios can be run in a simulation, the only things that usually can be controlled are the initial conditions and the occurrence of subsequent system or environmental changes. In complex systems, operator action is a significant variable in defining the path the scenario takes and by definition should remain uncontrolled. As a result, no two scenarios are exactly alike. Despite the fact that simulation offers the opportunity for comprehensive data collection, considerable work is involved in developing a useful performance measurement system for complex tasks. As a consequence, the studies that are performed tend to rely only on observational techniques and to produce performance descriptions that depend on the particulars

of the application, with few generalizable results (see Woods, 1984, for one exception).

Another potential disadvantage is that, unless experienced operators are available as test participants, the training of naive individuals can be a formidable task, and their performance may lack credibility. If one is interested in using simulation to test new control concepts, displays, etc., it requires extensive training on the part of the crew with a new system before meaningful data can be collected--training may be on the order of weeks, not hours or days (see the section below on subjects in supervisory control experiments). In many cases the amount of training required may be unrealistic, in practical terms of cost and time, to test, at best, a very small set of conditions. However, if an experiment can be integrated into training exercises conducted on high-fidelity simulators, the problems of training naive subjects can be avoided, provided the trainees are experienced operators undergoing annual or refresher training. This piggy-back approach has been used for research on nuclear power plant control (Woods et al., 1982). While high-fidelity simulation may have all the external appearances of an operational system, an important concern is whether the operators will react to simulated events in the same manner as they would to actual events.

Low-Fidelity Simulation

Low-fidelity simulation generally means partial or generic representation of a system that is useful for exploratory studies. Low-fidelity simulators offer considerable cost savings over high-fidelity, comprehensive simulations, while permitting somewhat more realistic and complex conditions, albeit for a narrow range of tasks, than is possible in laboratory studies involving only a few variables and highly abstract tasks. (In the future the difference between high- and low-fidelity simulators will decrease, as supervisory control rooms are reduced to a few display and control devices due to increased automation.) Low-fidelity simulation requires relatively limited amounts of practice by subjects, allows the observed behavior to be relatively easily analyzed and understood, and has modest data collection requirements.

Low-fidelity simulation, such as a downscaled plant, appears to be the most appropriate research vehicle for relating experimental psychological data to qualitative models based on artificial intelligence and control and decision theory; developing mathematical descriptions of human supervisory behavior; and distinguishing among skill-based, rule-based, and knowledge-based behavior. Experiments conducted at this level of simulation are likely to produce insights into the manner in which the human supervisor builds up an internal representation or mental model of the task to be performed, the system to be controlled, and the disturbance to be expected. While low-fidelity simulations reveal behaviors that, if present in the complex real system, would be of significance, it should not generally be assumed that behaviors will be the same in actual systems.

Studies at this level depend on simplifying assumptions about the task structure and environment. There are normative models and a large body of empirical results on decision making in static, discrete tasks, but, in actual systems operations, decision making is a dynamic process requiring coping with complex, changing environments (see Hogarth, 1981). In addition, actual systems consist of complex sets of processes in which functions are often controlled indirectly through the control of other processes (see Warfield, 1973). As a result, low-fidelity simulation can miss critical aspects of the subject-environment interaction. It should be viewed as a means for extending, elaborating, and evaluating laboratory findings and for identifying critical questions that should be addressed in the much more expensive, time-consuming, and complex studies that can be performed in high-fidelity simulators or operational settings.

Laboratory Studies

While research in actual operational settings and high-and-low fidelity simulations has the goal of understanding skill-, rule-, and knowledge-based behavior, it is often impossible to separate their contributions to the performance of complex tasks. Laboratory experiments are well-suited, at least for the study of general problem-solving strategies, to be used for knowledge-based behavior. Context-free tasks may also be a very good means for teaching knowledge-based behavior and to

gain insights about how the learning and experience gained on a conceptual level may be useful for real-world decision making and problem solving.

A promising approach for research on supervisory control behavior is to abstract critical features of the task environment (using some conceptual structure or theory of supervisory control behavior) into a generalized process model, and then to use experimental techniques to focus on the behavior of interest and thereby decrease reliance on verbal reports and naturalistic observation.

Laboratories are best suited for basic research on topics such as cognitive or mental models and the interactions between humans and abstract system representations incorporating artificial intelligence in the performance of complex tasks. It is also possible to test limited concepts, display devices, and workplace arrangements in a laboratory setting, and it is probably useful to do so. Small-scale laboratory experiments are appropriate for purposes such as evaluating ergonomic properties of display formatting and coding and matching specific display formats to selected decision strategies.

Even the simple "analogous task" has its place. It might be a computer game--computer games could be an effective vehicle for studying of human resources for direct symbol manipulation. (Willing subjects and data recording would certainly not be a problem.) The advantage of games is that they embody or illustrate certain new ideas and twists, serving as an operational definition of the construct. The disadvantage is that human subjects are apt not to take a game seriously.

Selection of Appropriate Research Vehicles

Useful experimentation on supervisory control systems is difficult, but research in real system, simulation, and laboratory settings all seems to have a place. Unless we begin with approximations, we will never develop the theoretical perspectives and methodologies necessary to understand supervisory control behavior and apply this knowledge to the design of systems in the future. The basic properties of supervisory control can be established initially using simpler systems for test and experiment, and the complexity of the controlled system can be gradually increased. Once the basic picture has been established using laboratory experiments, it will

be easier to decide on the critical experiments that can be carried out either on simulators or in real facilities. Although such initial laboratory experiments would be very economical, the use of full-scale, complex systems increases the cost of experiments considerably, in both the cost of running the real or simulated systems as well as in the cost of hiring or training experienced operators. In addition, only experienced operators will show "real" supervisory behavior.

The experimental plan or design, the research vehicle, and the subjects have to be chosen very carefully, since the appropriateness of each varies with the goals of the research. Table 2 summarizes the possibilities that should be considered in choosing a vehicle for research on supervisory control behavior.

Real-World Validation

For laboratory and simulation-based research it is important to recognize that there will always be differences in performance between simulators and real systems, and that laboratory and simulation results are not conclusive. No organization will or should base design recommendations strictly on laboratory results. Validation comes with research using real systems or at least full-scale simulators; only in full-scale simulation is the context adequate for the complex system interactions that we wish to generalize about. Operational evaluations are essential to see if the goals of the designer meet the goals of the user.

Thus, laboratory paradigms especially and low-fidelity simulation to some extent are devices only of intermediate usefulness. In fact, some would view the primary benefit of laboratory studies to be the training of researchers to have increased understanding of elements of the real task of interest.

Since changes to actual systems will be the inevitable consequence of operational testing, it would be highly desirable to design a system with a great deal of flexibility, install and use it operationally under careful monitoring and control, and, on the basis of operational testing, revise the system and only then release it for full operational use.

The Need for a Theory or Conceptual Framework

One of the most fundamental needs is for a conceptual framework or theory of supervisory control behavior to decide what aspects of the task environment are important to test, to guide or focus observational data collection techniques, and to provide a mechanism to facilitate integration of what is learned. It is impractical if not impossible to simulate all aspects of a system. It is therefore important to identify the critical features of the work environment with respect to the research objectives. These features may need high-fidelity simulation, while low fidelity may suffice for other aspects of the work environment. In other words, mixed fidelity may be necessary. Laboratory and context-free research may help develop the theory to identify which factors are critical enough to require high-fidelity simulation.

Standardized Forms of Research Vehicles and Tasks

It would be useful to arrive at some agreement among researchers as to what research vehicles and kinds of tasks are appropriate for studying the four control modes: start/stop procedures, normal process operation, fault management, and maintenance. For example, high-fidelity and even medium-fidelity aircraft trainers and regular flight simulators as well as nuclear power plant training simulators are appropriate for studying start/stop procedures, but laboratory tasks are not. Foxboro's computer-based simulator of a prototype chemical plant (Stassen, 1984) is appropriate to normal process operation. So is the "Tulga laboratory paradigm" (Tulga and Sheridan, 1980), wherein blocks of computer graphics appear and move at different speeds from left to right while the operator attempts to perform an exercise and earn a reward before the deadline. Reward is proportional to the area of each block completed, and completion time is proportional to width (Moray et al., 1983). Rouse (1980) has developed some simple laboratory tasks appropriate to fault management and maintenance, which involve tracing backward through a network of partially failed components to infer the cause of failure. A program packet for simulating a pressurized water reactor for supervisory control experiments is available from RISO (Goodstein, 1983) and

TABLE 2 Considerations in Choosing a Vehicle for Research on Supervisory Control Behavior

Research Vehicle					Games/ Context- Free
	Field Evaluation	High-Fidelity Simulation	Mixed-Fidelity Simulation	Laboratory Experiment	
Goals	Practical (positive as well as negative)	Practical/ Training	Tryout of conceptual ideas in reality	Conceptual	Conceptual training
Experimentation	Very costly; data collec- tion and processing difficult; performance assessment difficult	Costly; data processing difficult; performance assessment relatively difficult	Costly; data processing possible; performance assessment easy	Low cost; data collec- tion and pro- cessing easy; performance assessment easy	Low cost; data collec- tion and processing easy; performance assessment easy
Experimental design	Very difficult	Difficult	Easy	Easy	Easy

Results	Very practical; indicator for fundamental research	Practical; verification of new concepts	Verification of new concepts; model verification	New concepts and models	Training and new concepts, models
Transfer of research to design	High	High	Possible	Not direct	Probably direct
Process modes involved	All	All	[-----Difficult to do start/stop-----]		
Behavior involved	Skill, rule, and knowledge	Skill, rule, and knowledge	Knowledge	Skill, rule, and knowledge	Rule and knowledge

is scheduled to be used in 1984 in some major laboratories in Denmark, the Netherlands, the Federal Republic of Germany, the United Kingdom, and Canada.

SUBJECTS IN SUPERVISORY CONTROL STUDIES

The characteristics of subjects and the training they require are important considerations in the design of supervisory control studies. An essential element of experimental research on supervisory control behavior is to obtain valid performance, i.e., performance equivalent to that which would occur in actual operational settings under the same conditions. Often the investigator takes pains in selecting subjects to ensure that specialized learning or experience prior to that received in the experiment will not significantly influence their behavior. While one can probably use college sophomores to study basic perceptual and motor skills, it is very difficult to justify using such subjects in studies of decision making and problem solving.

Except for the most basic laboratory studies, in most investigations of supervisory control behavior the object is to understand decision and other cognitive processes that are largely dependent on extensive specialized learning. Consequently, to ensure valid performance, subjects in these studies must either be skilled operators or trained to an equivalent degree of skill. Experimenters should not place naive subjects in value-laden, decision-making situations that they have never seen before and will never see again and expect their decision making to reflect that of well-trained professionals who know that they will have to live with the consequences of their decisions. Of course, if the system of interest does not yet exist, experienced personnel will not be available. Nevertheless, analogous professional personnel are better than college sophomores.

The Training of Participants in Experiments

It is easy to underestimate the necessity for thorough training of subjects to be used in either full-scale simulation experiments or in experiments involving laboratory abstractions of tasks. Since the tasks are by definition complex and highly interactive, they

require both intellectual conceptualization and specific performance skills. It is known from both anecdotal information and formal research that the behavior of controllers changes over a very long time as a control skill is acquired. To build up prototypical control strategies, the trainee must experience a wide variety of conditions repeatedly. Stassen estimates that a minimum of 100-400 hours of experience is required in a typical case. Such a resource investment requires experimental designs that use small numbers of subjects serving in both experimental and control conditions in order to minimize the contribution of individual differences to measurement variability (see Towill, 1974, on the learning of time constants.)

The time required (time here is synonymous with number of trials, hours of practice, and extent of training) is a function of system complexity. For simple monitoring tasks with systems having one or two degrees of freedom, a level of skill that is close enough to asymptotic for useful analysis can be obtained in around 10-20 hours. For full-scale industrial control skills in a complex plant, manual control skills improve over hundreds or thousands of hours. For the particular kinds of skills involved in supervisory control, in which little manual control is exercised but it is important for the operator to understand the system, the time to quasi-asymptote will be in the higher ranges. The time required will be a monotonically increasing function of system degrees of freedom, a monotonically inverse function of the system bandwidth (assuming the bandwidth is less than 1.0 Hz), and a monotonically increasing function of the number of levels in a hierarchical system. The effect of the amount of coupling between subsystems is unknown.

If the researcher has good reason to believe that an operator has settled permanently on one strategy of operation, and that practice is merely increasing the efficiency with which he or she is exercising that strategy, then the amount of practice required before useful performance data can be collected will be reduced, since skill acquisition occurs on a smooth curve that is either a log-linear or a log-log function of practice. From the shape of such learning curves it follows that most learning takes place during the first few tens of hours, even though skill will continue to improve at a slower rate over hundreds of hours. This consideration means that the problem of training to

criterion is not as serious as it may at first seem. Subjects doing very simple tasks with two or three degrees of freedom will often show most learning within five hours or less (Moray et al., 1983).

If a number of acceptable alternative strategies exist, however, discontinuities in the learning curve may occur, and after each such discontinuity it should be assumed that a further substantial period of learning, at least on the order of tens of hours and perhaps longer for complex systems, will be necessary.

In many cases there is probably more than one strategy available to operators to bring a system to a particular desired state. It is also probable that an operator who discovers such a strategy early in practice will continue to use it and will become very proficient at it, but will not learn system properties that are related to alternative strategies.

Data Collection During Training

On one hand, every opportunity should be taken to collect data during the training phases of experiments, especially for full-scale simulator or field experiments. On the other hand, it seems likely that periodic sampling will be sufficient to give the overall picture and enough details to follow the process of skill acquisition. Data on eye movements as well as control manipulations should be collected whenever possible. This is especially important as practice continues, since the number of control manipulations per hour decreases markedly with skill acquisition.

Besides raw data and learning curves, transition matrices for every response contingent on every preceding response should be analyzed (link analysis), since this can give very interesting insights (see Moran, 1983, for example). Error and latency should be measured. It is also possible that information analysis as developed by Conant (1972) may be useful for the analysis of structure.

In addition to collecting data during normal operations, data recording behavior during fault detection and diagnosis should be gathered.

Individual Styles of Performance

A complicating factor in the choice of subjects is that regardless of whether two people have had identical training, their performance on the same tasks may differ substantially. There is conclusive evidence that, even for relatively simple tasks, different operators use different styles of control. There are several dimensions on which operators may differ, but experimental evidence exists for at least the following dimensions for a variety of tasks:

1. Emphasis on speed versus emphasis on accuracy when making decisions.
2. Emphasis on immediate detail versus emphasis on overall context in acquiring and interpreting information (field dependence versus field independence).
3. Emphasis on the present context versus the use of past experience in relation to the time span over which extrapolation or prediction occurs.
4. Cognitive fixedness versus cognitive flexibility.

Although much psychological literature implies that such differences in individual styles ("cognitive styles") are caused by stable differences in personality, there is empirical evidence for major changes as a result of training and the exercise of skills. (For example, experienced radar operators are more field independent than beginners, but become less so when they cease to actively perform fighter control tasks--Moray et al., 1982). Such dimensions as speed versus accuracy are sensitive to perceived event probabilities and perceived payoffs, which in turn may be affected by peer pressure, training, implicit or explicit management policies, etc.

Differences in individual style will affect the speed and accuracy with which an operator completes a task, the amount and kind of information he or she requires before making a decision, and his or her proneness to cognitive fixedness. The design of displays, controls, and data entry devices will interact with the operator's style, and conflicts of style between members of a team may be advantageous (in allowing a variety of approaches to a task) or disadvantageous (by making it difficult for members of a team to agree on criteria, diagnoses, and actions).

Implications for Research

Differences in subjects' experience and amount of training have several practical implications for research:

1. The continuum shown in Figure 4 can be thought of as a continuum from "relatively little practice required" to "very extensive practice required."
2. In terms of cost effectiveness, issues involving experimental subjects should be clarified using settings at the low end of the continuum and validated at the high end.
3. Well-qualified, highly trained subjects are required for supervisory control research. Selection of the subjects may have an impact on the research. In the case of professionals who are "borrowed" from their usual tasks, there will be positive transfer of skill from the normal task to the research task to the extent that the two are similar (i.e., including payoff structure, management policy, physical appearance of the control room, process dynamics, etc.). There will be negative transfer to the extent that there are differences. Training new operators is costly and time-consuming.
4. Frequently the researcher is interested in comparing a new design or system structure with an old one. Under these conditions there is a real danger that the subject's experience with the old conditions will interfere with effective utilization of the new conditions, the problem of negative transfer of training. These issues taken together are very difficult to resolve using existing formal experimental designs. It may be possible to limit the scope of experiments so that the training requirement itself is reduced. Methodological developments addressing this paradox are urgently needed.

PERFORMANCE ASSESSMENT

While the measures and criteria depend on the specific purpose of each investigation, the general purpose of all performance assessment is to obtain information for decision making by researchers, designers, managers, and others.




FIGURE 4 Progressive Complexity of Research Vehicles

Just as supervisory control implies a hierarchy of levels of control, so too does it imply a hierarchy of assessment measures and criteria. Traditional measures, such as response latency and root mean square error, are most suited to the lowest levels of control (except perhaps for responses to emergencies). At management levels, measures such as productivity or efficiency may best reflect the controller's behavior. At the social level, net earnings, rate of earnings, peer approval, or job satisfaction may be appropriate, since these are the personal objectives the controllers may ultimately be striving for.

Sometimes the performance criteria at any of these levels must be seen as a trade-off, e.g., keeping a plant running versus taking safe actions, or landing an aircraft on time versus doing it safely. The "control philosophy" is intimately related to assessment criteria, as are labor relations, safety, manning levels, etc.

Assessment of the performance of operators of supervisory control systems is complicated by the facts that these systems require relatively few overt actions by the operator, and the significance of an action at any given time may vary considerably. In general it is the operator's decision processes that result in certain actions or inaction that are of interest; more specifically, it is the factors that influence these decisions, e.g., the state of the system variables, what information is available, preceding events, the operator's knowledge and mental model of the system.

One of the most important concerns in the study of supervisory control behavior is error resulting from incorrect decisions. A comprehensive, in-depth analysis of the situational context is probably the most critical factor in the assessment of the causes of decision errors. Actual measurement of overt actions is a somewhat minor problem compared with constructing a framework, and hence criteria, for interpreting the error actions of the operator. While there are

documented cases in which unexplainable, erratic behavior produces system errors, it is far more common for the controller to take actions for reasons that appeared rational from his or her perspective.

In-depth analysis requires that the full context in which the control action occurred be understood, not just the outcome. This can be achieved only if the actual event or the simulator scenario is carefully analyzed in terms of the structure of decision alternatives it presents to the controller. Given this context it is then possible to examine each operator decision, solicit from the individual an understanding of the context at the time, including expectations about the consequences of particular control actions. From such an examination an entirely new set of performance measures will emerge that are much more likely to reveal insights about how the system is being used and how to structure design alternatives. It is an analysis based on decision making rather than simply on outcomes. Such analysis is also useful as the basis for evaluating controller performance in a training situation. Initial attempts to structure this kind of analysis have been made (see Pew et al., 1981; Hollnagel et al., 1981; Woods et al., 1982); these experiences need to be expanded and standardized so that data from a wide variety of contexts can be used to understand controller performance in more general terms.

Some particular decision-making problems occur frequently in a variety of contexts. One of these is that decision errors tend to persist. Research should be targeted specifically on discovering the behavioral mechanisms that underly these fixation-like, "cognitive lock-up" phenomena. At another level, research is needed to develop system architectures that prevent error, ameliorate persistence of error, and facilitate error correction. This approach might be characterized as "cognitive defense in depth."

Methodologically, decision analysis (e.g., Raiffa, 1968; Fischhoff, 1980) might appropriately be applied to operation of supervisory control systems. In accident investigation, decision analysis has proven to be a useful way of reconstructing how operators viewed their situation and where they went wrong in coping with it. The approach has the advantage of being nonevaluative. It assumes that operators' actions were purposeful and well-intentioned. It recognizes whatever inherent uncertainty a system has (given what an operator can do

or does know about its operation). It allows actions to be seen as gambles, and bad luck to be a potential cause of accidents. (Indeed, acknowledging uncertainty allows some calculation of the amount of bad luck to be expected, such calculations being the province of risk analysis). Once the general locus of difficulties has been identified, more detailed investigations can be conducted regarding the areas in which redesign or retraining are needed.

4: THE DESIGN OF SUPERVISORY CONTROL SYSTEMS

SPECIAL DESIGN CONSIDERATIONS

Modern systems are increasingly large, complex, semiautomated, and under centralized supervisory control. While automation may relieve human operators of direct, detailed, and continuous involvement in process tasks, it places increased demands and responsibility on the human supervisor for monitoring the function of the overall system, detecting out-of-tolerance conditions and faults, and deciding on and executing timely, correct actions. Providing the appropriate means and mechanisms for a system supervisor to perform these tasks is a critical aspect of the design of supervisory control systems. Inadequate or inappropriate displays and controls can have costly and dangerous consequences.

Design of supervisory control systems requires a somewhat different perspective on the role of people as system components from what was appropriate in the past. The principal role of the human supervisor is to be a decision maker whose actions are an input to an otherwise autonomous, closed-loop process. In older systems the person was an active control element in the process loop with actions required moment to moment. In older systems the equipment functions were responsive solely to the control of the human operator in a fixed, deterministic way. New systems incorporating machine intelligence do not necessarily act in a predetermined manner as a consequence of the supervisor's inputs. That is, these systems have the capability, within

limits, to determine what means are best suited to achieve the goal established by the supervisor's actions.

The Goal-Directed Character of Supervisory Control

Designers of supervisory control systems should recognize explicitly the goal-directed character of supervisory control: a system demonstrates purposive behavior if it can adjust itself in the face of changing conditions in pursuit of a goal. Typically, there is a complex goal topology, i.e., a many-to-many mapping between resources (means) and goals, wherein a given means may influence several goals, and goals are potentially supported through many means. This analysis suggests that supervisory control activities consist of the interaction between goal-directed search (Given an interest in achieving a goal, what mechanisms support that goal?) and goal-oriented assessment (What are the implications for achieving a goal of the various resources available?). In this view, hierarchical models of human mental processes, particularly decision processes, can serve as a basis for modeling the human-machine cognitive system in supervisory control.

The Hierarchical Structure of Supervisory Control Systems

Given that supervisory control systems are highly automated and therefore can respond appropriately to routine or anticipated perturbations, the principal task of the human supervisor, once the system process is started, is to make decisions and take actions when rare, unusual, or unforeseen perturbations or failures occur. Therefore, a design goal should be to provide the supervisor with mechanisms to gain the appropriate information necessary for making decisions and to exercise control at any level in the system without the designer's knowing specifically what events may occur, what information the supervisor may need, or what actions will be taken.

To support the supervisor, the information and control resources must allow for reconfiguration of the system in terms of the basic purpose, function, and equipment relationships at different levels of the system. Therefore, in designing the information access and control mechanisms, an essential requirement is a consistent hierarchical representation of these relation-

ships. Such a representation provides a framework of the context in which decisions are made. This is a prerequisite for developing a decision-making model-- "The complexity of human behavior largely reflects the complexity of the environment" (Simon, 1969). The hierarchical representation is also necessary for evaluating the adequacy of resources to meet the supervisor's needs for coping with perturbations or failures. That is, it is necessary for systematically identifying constraints on decision making and possible deficiencies, e.g., "sneak-paths" or unforeseen but possible sets of conditions that may lead to decision errors. Furthermore, a hierarchical representation can aid in developing the nature of supervisory control tasks, procedure development, and maintenance planning.

Human Performance, System, and Decision Models

The successful design of control rooms for supervisory control systems depends on the availability of: (1) models of human performance for determining cognitive task allocation; (2) models of system properties and decision processes to identify control task requirements; and (3) the design of the supervisor-to-system interface. These resources were largely unnecessary when interface design was based on the traditional one-sensor, one-indicator and fixed-function control technology. The central importance of cognitive and decision process models to the design of future supervisory control systems implies not only the need for substantial research to fill the designers' needs, but also that system designers will have to change the way they design.

The Interface Between the Supervisor and the Intelligent Computer System

The use of intelligent computer systems in supervisory control creates a new generation of interface design questions--not only person to process, but also human cognitive system to machine cognitive system. For example, in certain robotic systems, the problem initially is to teach the robot what to do, which in turn (given some intelligence on the computer's part) requires some context to the problem, i.e., initial conditions to be set in a way that both the person and the computer understand.

Design questions include issues of authority, responsibility (even legal liability), and locus of control (e.g., Fitter and Sime, 1980; Woods, 1982):

- o Who is responsible for system performance--the operator, the computer, or the decision system designers?
- o Should the operator at times be a servant commanded by the computer to follow a decision or take an action?
- o Should person and computer work as equals who cooperate to reach a solution?
- o Or is the operator, as the responsible party, in charge, with the intelligent computer system serving as another resource to be managed to achieve process goals?

One example of problems with cognitive system interface occurs when a system produces an answer to a given problem. Regardless of the form of the suggested answer (a command, a probability, or "advise"), there is great danger of its constraining the human supervisor to either reject or override the system output (perhaps by finding or creating grounds for system unreliability) or to abrogate the decision responsibility (i.e., regardless of circumstances, the operator may not override the computer, if the cost of an error in overriding is too high). This double bind is exacerbated if the operator cannot determine the basis for the computer's answer.

Another example occurs when an intelligent computer system controls the process the person is used only as a resource to supply data or take actions when directed. When the locus of supervisory control resides with the computer, there may be no person-computer synergy. The operator may have, as we know from experience (British Steel Corporation, 1976), great difficulty assessing the achievement of process goals, identifying problems in the performance of the automatic decision system (if a failure occurs or if a process disturbance is too difficult for the decision system to handle satisfactorily), and adjusting or taking over the decision system function when necessary.

These examples point out that if people are to be responsible for the outcome of their own decisions, the intelligent computer may need to be designed as another resource they use to perform supervisory control. This suggests that computer intelligence should be used less

to make or recommend solutions and more to aid supervisors in the process of reaching a decision.

The success of the computer in aiding the decision process will also depend on the willingness of operators to adopt decision strategies for which the computer intelligence is designed. There is a delicate balance between designing computer intelligence to complement and aid a broad range of styles of decision making and training people to adopt particular strategies to use the computer resources most effectively.

This last point emphasizes the need for psychological and human performance models of supervisory control behavior that represent the cognitive processes, especially decision strategies, used for particular classes of problems. These models should reflect the supervisor's subjective goal formulation and choice of decision strategies as they would occur in the actual operational setting. Such models would be helpful to determine the supervisor's acceptance and use of computer support. This knowledge, in turn, could be applied in the design of intelligent systems to aid human decision making and the formulation of design policy for the allocation of tasks between the human supervisor and the intelligent computer. These models would also be useful in the development of training in decision strategies.

A simple model is needed of the moment-to-moment values to the operator of various decisions or actions in order to develop a policy for this allocation of tasks, as well as a calculation of how the operator will respond. Manipulation of the payoff structure for the operator may influence the extent to which he or she calls on the computer for assistance or control (Moray, 1981).

Staffing Levels and Policies

Part of the analytic process of design must be to decide on staffing levels, which must be an iterative process. The management or the designers propose a staffing level. This must be related to the prevalent philosophy of control. Task analysis and knowledge of human operator characteristics will then be used to specify how proposed staffing relates to control room design in terms of displays and controls. If the results show that the level of activity required is too great for the staffing level proposed, then there must be recommendations for a

change in these levels. If adopted, this will change the design of the control room procedures, etc. This process then should be repeated. The opinions and desires of the operators as well as designers should be part of the process, a factor that argues for a mock-up and simulation early in the design phase.

COMMUNICATION BETWEEN RESEARCHERS AND DESIGNERS

Supervisory control, because it is so encompassing, poses great demands on communication between researchers and designers. Researchers can help designers by contributing facts, principles, and design concepts that serve as inputs to their process of synthesis and analysis; they can also outline the inevitable trade-offs among competing requirements for human abilities and the possibilities of exceeding human limitations. These types of information can be supplied in the form of lectures, textbooks, handbooks, and guidelines as well as participation in design efforts. Success in transferring knowledge from research to design, using any of these modes of communication, is heavily dependent on the degree to which the timing and form of the information complements the design process. Tabulations of potential answers to human factors design questions are of little use if the designer does not realize, at the appropriate time in the design process, that he or she should be considering these questions.

For research to be conducted and disseminated in a way that will be useful to designers, there has to be greater mutual understanding between designers and researchers. The direct way to achieve such understanding is through personal contact at professional meetings, conferences, and working groups. The great advantage of such contacts is that they allow conversation, the exchange of views, and requests for clarification. The primary disadvantages are that the contacts are necessarily selective and the process can be inefficient and expensive. Only a sampling of researchers or designers get to such meetings and working groups, and they may not be representative of their professions as a whole. Perhaps more important, only a sampling of a researcher's or designer's working life is mutually observable in these contexts (and it too may be unrepresentative of the whole). As a result, even such contacts may leave an incomplete or inaccurate

picture. A partial solution to this problem, at least as it confronts researchers, is to conduct systematic studies of the design process.

Publication of journal articles, conference papers, and technical reports is reasonably efficient and inexpensive, but the effectiveness of such communication is questionable. In archival journals, researchers publish articles to be read by other researchers. Conference papers and technical reports may be more readable from a designer's point of view, but the quality is very uneven and the distribution usually quite limited. Design handbooks and guides would seem to be a reasonable vehicle for communication, but unfortunately these tend to be compendia of facts rather than aids to the process of design. This situation is at least in part due to the fact that design is a poorly understood process.

Design Resources

The introduction of complex systems into an operational setting has produced some unanticipated side effects in terms of human factors. This is a result of our current lack of understanding of the people's reaction and adaptation to these systems. Examples of such behaviors are the bypassing of safety systems, the loss of skills, and additional training requirements. Because these responses may be unanticipated, they have the potential of leading to events that are dangerous or costly. Even if designers can anticipate some of the side effects, they may not have enough hard data to argue for design changes that have significant cost consequences. Therefore, the best opportunity for influencing a system's design is in the early stages, with the prospect of turning out a better design on successive iterations. To do this designers must be provided with information about human factors principles relevant to the design of supervisory control systems (Stassen, 1984).

Principles are probably the most effective way to convey to the designer the necessary information for incorporating human factors considerations in the design and operation of supervisory control systems. Principles are those findings that are robust (e.g., frequent false alarm will tend to be discounted or ignored), even though the parameters and levels may not be known in any

single situation. The assumptions and background must be stated for each.

Past experience has shown that, when principles are used by designers, they can result in well-designed equipment from the human factors point of view. Principles have several advantages. They are intuitively appealing, and they can be understood and incorporated by the designer in the early stages of system design, when there is little or no penalty for their application. One particular value is that they can help the designer look for shortcomings in the design. Principles can also provide a common basis for discussion and design review. They provide a means for accumulating knowledge within an organization. Capitalizing on past experience, they can be expanded, become more specific, and be tailored to meet the needs of a particular organization.

As an adjunct to principles, checklists should be provided to help the designer identify the critical items and behaviors that may influence system performance.

To apply principles properly a designer must also be made sensitive to the issues that determine the applicability of guidelines in specific contexts and to the questions that have different answers in different situations (e.g., is failure detection by the operator better when he or she is a controller or a monitor?). Each issue must be explained to allow the designer to draw the correct conclusion for each particular case.

In summary, information from psychology, human-machine systems, and operational experience should be translated into principles, checklists, and explanations of issues to aid the designer and should be in a form that can be easily modified and updated to meet the needs of different organizations.

Accumulating and Sharing Knowledge

There is a need for more emphasis on the accumulation of lessons learned from actual experience with both successful and unsuccessful designs and the sharing of this information among organizations working on similar problems.

While research is being conducted on the theory and concepts of supervisory control, a great deal of useful information about current systems exists, but is not known widely enough. The one thing that remains

relatively constant throughout the computer revolution is human nature. A concerted effort must be made to collect the many lessons learned both within and across industries: this could take place at many levels, from dial or gauge design through systems architectures and information management.

In addition, negative results and knowledge gained from unsuccessful designs are as important as positive results and successful designs at this stage of our understanding of the people's interaction with complex systems. Such information is often hard to get and must be coaxed out in a form that is anonymous in terms of the organizations or institutions.

Ideally, mechanisms for research and dissemination of information on supervisory control should include cooperative international efforts. Such efforts might include a laboratory for supervisory control research, which would involve a real process and a full-time cadre of employees, funded as an international project between the United States, Canada, and other countries, along the lines of the Haldane laboratory for nuclear power research in Haldane, Norway, a data bank for general use, and an international information exchange network to supplement traditional journals and meetings.

Guidance From Research

Designers need guidance on the system features required to support performance of the supervisor's tasks. This guidance necessarily must come from research on the nature of supervisory control tasks and the behavioral processes involved in their execution. Some of the more important research questions include:

- o What are the requirements for supervisory control displays and interfaces that characterize the state of the control of the process, not just information on the state of the process (Woods, 1982)?
- o Errors of decision making tend to persist (e.g., Woods, 1984); how can system designers avoid these fixation-like effects?
- o How can the supervisor develop or maintain the control skills needed when automatic systems fail (Bainbridge, 1982)?

- o How, in "teaching" the supervisory control system a new task, can the supervisor know what the intelligent system already understands (Sheridan, 1984a)?
- o How do characteristics of function allocation and interface affect the operator's ability to develop an effective model of system dynamics, including, for example, learning dynamic target states, models of slowly responding systems, and process feel (Ephrath and Young, 1981; Wickens and Kessel, 1981)?
- o To what extent can people make use of correlational information from several variables that are tightly coupled (Moray, 1981)?
- o What are the design requirements for operational systems in which several levels of supervisory control are distributed over people, facilities, and intelligent computer systems (Sheridan, 1984a)?
- o How can breakdowns in supervisory control be prevented or minimized, which implies a model of human supervisory control performance (North Atlantic Treaty Organization, 1983)?
- o How does supervisory control performance depend on the coordination of behavior as distinct from separate specific activities of operators?
- o How can we design the interface of the intelligent computer system with the human supervisor to maximize system performance (Hollnagel and Woods, 1983)?

Understanding the Design Process

Design is sometimes viewed as a much more analytical process than it actually is. Designs emerge from an iterative synthesis-analysis process: while evaluation tends to be more analytical, synthesis is known to be fairly chaotic. For research results to be useful, they must be integrated into this process at appropriate places. As a practical matter, such integration requires a better understanding of the design process than seems to exist currently.

An experimental approach to understanding the design process for supervisory control systems might have designers predict how various aspects of their designs will actually perform, how operators will use them in

particular situations, or what degree of trust (or acceptance) they will receive. A more sociological (or anthropological) approach would be to observe a variety of actual design processes (or reconstruct them after the fact). Through interviews and observation, a picture would be created of how designs evolve, where ideas come from, how explicitly goals are stated, how (or if) conflicts between goals are reconciled, whether designers themselves can reconstruct the process, when (or if) human factors experts are consulted, when (or if) operators are consulted, how much (if any) flexibility is left to change the designer's response to experience, and how much latitude designers have to accept new ideas. Even if there were no substitute for direct contact, some good case studies could make those contacts easier, by showing how designers think, speak, and work. They could also provide a partial answer to the crucial question of whether the supervisory control perspective can be sold as a whole or whether the best that can be hoped for is the diffusing of a few principles that would improve, without revolutionizing, the design process.

To ensure that the right questions get asked at the right times, a better understanding of the process of designing a supervisory control system is necessary. Based on this understanding, research information could be transformed to match design needs. At first this will take the form of new types of handbooks and guidelines. Later, it may be used to transform design education, where a substantial and lasting transformation is most likely to have the greatest effect.

5: CONCLUSIONS

This final chapter presents conclusions and their implications for research that emerged from discussions at the workshop. They draw on the three main themes that structure this report: (1) characteristics and analysis of supervisory control systems, (2) selection of research methods, and (3) improved communication between research investigators and designers of these systems.

CHARACTERISTICS OF SUPERVISORY CONTROL

Levels of Control. The trend in supervisory control systems is toward two levels of control with two computers, each having different functions. The human supervisor's interaction with a computer may be called a human interactive subsystem (HIS) and control of the final product or process the task interactive subsystem (TIS). There is great flexibility in the couplings among the human operator, the HIS, the TIS, and tasks, with a fanning out or multiplexing of control at the lower levels. The HIS can be more or less intelligent in terms of giving advice to and helping the operator plan ahead and learn from experience.

Goal Setting and Seeking. Supervisory control systems are hierarchical in nature, and the hierarchy can be described in terms of goal setting and seeking. Goals and specifications propagate from the top down. Resources and limitations are supplied from the bottom up. Control tasks (what is to be done) should be formulated for the process at each level.

Cycles of Control and Feedback. In supervisory control there are typically five-steps of supervisory behavior: (1) planning, (2) instructing the computer, (3) monitoring and adjusting, (4) intervening to re-program or to override automatic control as necessary, and (5) learning from experience. Iterative feedback usually occurs between learning and planning at long intervals, between intervening and instructing the computer at shorter intervals, and, at very short intervals within the monitoring and adjusting step.

Trust. The operator's perception of the trustworthiness of a system affects his or her behavior. For example, without a sense of assurance that the commands will be carried out properly at a lower level, the operator may spend considerable time checking for compliance or may simply bypass automated control. The appropriate criteria for trusting should be studied to develop a theory of trust in supervisory control.

ANALYSIS AND MODELING OF SUPERVISORY CONTROL SYSTEMS

Methods for analyzing and modeling supervisory control systems are necessary to understand and predict the supervisor's behavior, to provide a common, formal basis for comparisons among systems, to allocate tasks between the supervisor and the computer, and to differentiate the contributions of the system's hardware, software, and personnel resources.

Modeling. Both qualitative and quantitative models should be developed. The designer in effect must make explicit a model for achieving the purposes the system is intended to fulfill. The system may need a model of itself to achieve good automatic control as well as some model of the human operator to help the two communicate. The operator has a mental model of the task or process and a model of the automatic controller. A goal for design is to have all these models in harmony.

Human Error. System failures resulting from the operator's decisions or actions are often viewed as the result of a mismatch between the person and the system. Such human errors in supervisory control systems certainly include slips, forgetting, and other common types. Another type of failure results from the operator's exercising control too frequently or too infrequently (due to either mistrusting or overly trusting the system). Preparing the operator to cope

with unexpected events and to avoid errors is one of the central problems of research on supervisory control systems.

VEHICLES FOR RESEARCH

Research to understand supervisory control behavior is undertaken to develop behavioral principles as a basis for design and to evaluate design alternatives. The research can be supported by real systems, high- and low-fidelity simulation, and context-free laboratory tasks.

Real Systems. Real systems are the source of the problem set and the final proving ground for the design, training, and procedural concepts developed in other research contexts. Operational systems are the best vehicle for the analysis of such things as the subjective goals of the operator, information processing strategies actually used as a function of experience and training, and trust in the system's control mechanisms.

Real systems present several difficulties as research vehicles. They span a wide range of physical characteristics, operating philosophies, and training procedures. Behavioral data on rarely occurring, abnormal events are difficult to collect and are likely to be derived by observation or eliciting verbal reports from operators.

Simulators. Simulators have several advantages as research vehicles. Conditions and events can be created and controlled as necessary, detailed data can be collected, and often several experimental questions can be resolved in a single study.

High-fidelity, comprehensive simulation can be used to study rarely occurring events in a real-world configuration. Low-fidelity simulation attenuates both the advantages and disadvantages of high-fidelity simulation. Low-fidelity simulation appears to be most appropriate for relating fundamental psychological data to qualitative models and for developing mathematical descriptions of components of supervisory control behavior.

Laboratory Settings. Laboratories, where some modicum of experimental control is possible, are best suited for basic research on cognitive processes and interactions between humans and abstract system representations incorporating artificial intelligence in

the performance of complex tasks. The laboratory is also an appropriate setting for evaluating display properties.

The Importance of Validation. Economics sometimes dictates using simpler research vehicles for initial test and experiment, and gradually increasing the complexity of the representations of the system under study. However, real-world evaluation or at least full-scale simulation may ultimately be necessary to validate principles of supervisory control behavior. Only in this context do the complex interactions occur that are essential to confirm that the goals of the designer and the goals of the operator are the same.

Standardized Forms of Research Vehicles and Tasks. It would be useful if researchers could agree on what research vehicles and kinds of tasks are appropriate for studying each of the four control modes, i.e., start/stop procedures, normal process operation, fault management, and maintenance. Some commonality of research vehicles across laboratories would greatly facilitate cooperative efforts and comparison of results.

SUBJECTS, THEIR STYLES, AND PERFORMANCE

Subject Requirements and Training Requirements. In most investigations of supervisory control behavior the object is to understand decision and other cognitive processes that are largely dependent on extensive specialized learning. Consequently, to ensure valid performance, subjects in these studies must either be experienced operators or trained to an equivalent level of skill.

Frequently the researcher is interested in comparing a new design of system structure with an old one. In that case there is a real danger that the experience of subjects with the old version of the system will interfere with their effective operation of the new one. This paradox--requiring experienced subjects but having the experience hinder performance--is very difficult to resolve using existing formal experimental designs. Methodological developments to resolve this paradox are urgently needed.

Training. The time required to train subjects varies enormously: it can be quite short for simple laboratory tasks, and it can involve hundreds of hours if the tasks are comprehensive representations of large-scale, complex, industrial control plants.

The number of possible alternative strategies that must be learned has a greater effect on training time than does achieving some criterion of performance with a single strategy. The process of learning to exercise supervisory control is an important research topic.

Individual Control Styles. A complicating factor in supervisory control experimentation is that each operator exhibits individual styles of control for even relatively simple tasks, e.g., emphasis on speed versus accuracy, on immediate detail versus the overall context in acquiring information, using the immediate context versus past experience to solve problems. These differences, although they complicate performance measurement and analysis of experimental results, are a reflection of reality and should be taken into account in the research plan.

Behavioral Measures. Supervisory control involves hierarchies of goals and levels of control. Traditional measures of performance such as response latency and root mean square error are most suited to the lowest levels of control. At the highest supervisory level, measures such as productivity or efficiency may best reflect the operator's behavior. In between, measures of information seeking, decision making, and control strategy are appropriate.

One of the most important issues in the study of supervisory control behavior is that of error resulting from incorrect decisions. Promising attempts to analyze and structure the decision alternatives available to an operator need to be expanded and standardized so that the data from a wide variety of contexts can be used to assess and understand performance in more general terms.

DESIGN OF FLEXIBLE OPERATOR-SYSTEM INTERFACES

The operator must have display and control mechanisms to acquire the information necessary to make decisions and to exercise control at any level in the system. In designing the system, however, the designer will not know what information may be needed or what actions may be taken.

Representing Display and Control Requirements. An essential for designing effective mechanisms of information access and control is the development of a consistent hierarchical representation of the purpose, functions, and equipment relationships at different

levels in the system. This representation would describe the context in which operator decisions are made as well as provide a basis for evaluating the adequacy of resources to meet the operator's needs in coping with abnormal conditions.

The Interface Between Operator and Intelligent Computer. A new generation of interface design questions centers on the interrelation of human cognition and machine cognition. These questions include issues of authority, responsibility and locus of control as well as compatibility between the human and computer representations of the system.

RESEARCH ISSUES RELEVANT TO DESIGN

Specific design considerations include whether it is better to display the state of control rather than the process state; the value of displaying correlational information from tightly coupled variables; prevention of the persistence of decision errors; how the allocation of tasks between the supervisor and the computer affects the supervisor's ability to develop an effective mental model of the system dynamics, particularly for slowly responding processes, and how the operator might develop and maintain manual control skills necessary when automatic control fails.

COMMUNICATION BETWEEN RESEARCHERS AND DESIGNERS

Success in transferring behavioral research data, principles, and concepts to designers depends on the extent to which the content, timing, and form of the information is compatible with the design process and on awareness on the part of the designer of what questions to consider. How to improve the use of research information and promote better communication between designers and researchers is not yet clear.

Personal Contact. Personal contact between researchers and designers at professional meetings is the most direct method, but it is selective, limited, and expensive. Journal articles and technical reports are efficient and inexpensive but their effectiveness is dubious; they are more likely to be read by other researchers than by designers. Behaviorally oriented design handbooks seem to be the most reasonable means of

aid design rather than being only compendia of facts.

Design Resources. Human factors principles and checklists, accompanied by statements of the issues that determine their applicability, can be very helpful to supervisory control systems designers.

Principles have the advantages of being intuitively appealing, easily understood, and useful in the early stages of design when there is no penalty associated with their application. They have the disadvantages of being vague and nonquantitative and not providing information on trade-offs among conflicting design objectives. Checklists as adjuncts to the principles can help the designer identify the critical items and behaviors that may influence system performance.

Accumulating and Sharing Knowledge. An important source of design information is lessons learned from existing systems. There is sufficient similarity among different forms of supervisory control systems to warrant efforts to promote sharing of experiences across industries. Negative results and knowledge gained from unsuccessful designs are as important as positive results, given the current state of understanding of supervisory control behavior.

Understanding the Design Process. A better understanding of the design process is needed to make behavioral information useful to supervisory control system design. The design process may be very chaotic: designs emerge from an iterative cycle of synthesis and analysis. Efforts to understand the design process could provide at least a partial answer to the question of how influential the behavioral considerations are likely to be in the design of supervisory control systems and whether human factors principles can make a substantial contribution to their improvement.

REFERENCES

- Bainbridge, L.
1982 Ironies of automation. In G. Johannsen and J.E. Rijnsdorp, eds., Analysis, Design, and Evaluation of Man-Machine Systems. New York: Pergamon Press.
- Baron, S., Seehrer, C.E., Rew, R.W., and Horowitz, P.
1982 An Approach to Modeling Supervisory Control of a Nuclear Power Plant. Report no. 4811. Bolt Beranek & Newman, Inc., Cambridge, Mass.
- Boff, K., Kaufman, L., and Thomas, J., eds.
1984 Handbook of Perception and Human Performance. New York: John Wiley and Sons.
- British Steel Corporation
1976 Human Factors Evaluation, Hoogovens No. 2 Hot Strip Mill. London: British Steel Corporation.
- Conant, R.
1972 Detecting subsystems of a complex system. IEEE Transactions on Systems, Man and Cybernetics. SMC-2:550-553.
- Dawes, R.M.
1979 The robust beauty of improper linear models in decision making. American Psychologist 34:571-582.
- Ephrath, A.R., and Young, L.R.
1981 Monitoring vs. man-in-the-loop detection of aircraft control failures. In J. Rasmussen and W.B. Rouse, eds., Human Detection and Diagnosis of System Failure. New York: Plenum Press.

- Ericsson, K.A., and Simon, H.A.
 1984 Protocol Analysis: Verbal Reports as Data. Cambridge, Mass.: MIT Press.
- Ferrell, W.R., and Sheridan, T.B.
 1967 Supervisory control of remote manipulation. IEEE Spectrum 4:81-88.
- Fischhoff, B.
 1980 Clinical decision analysis. Operations Research 28:28-43.
 1982 Debiasing. In D. Kahneman, P. Slovic, and A. Tversky, eds., Judgment Under Uncertainty: Heuristics and Biases. New York: Cambridge University Press.
- Fitter, M.J., and Sime, M.E.
 1980 Responsibility and shared decision-making. In H.T. Smith and T.R.G. Green, eds., Human Interaction With Computers. London: Academic Press.
- Gentner, D., and Stevens, A.L.
 1983 Mental Models. Hillsdale, N.J.: Erlbaum.
- Goodstein, L.P.
 1982 An integrated display set for process operators. Pp. 75-82 in Proceedings of IFAC/IFIP/IFORS/IEA Conference on Analysis, Design, and Evaluation of Man-Machine Systems.
 1983 The GUP Test Bed for Operator Support Evaluation. RISO-N-16. RISO National Laboratory, Roskilde, Denmark.
- Hammond, K.R., Hursh, C. and Todd, F.J.
 1964 Analyzing the components of clinical judgment. Psychological Review 71:438-456.
- Hogarth, R.M.
 1981 Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin 90:197-217.
- Hollnagel, E., and Woods, D.
 1983 Cognitive systems engineering: New wine in new bottles. International Journal of Man-Machine Studies 18:583-600.
- Hollnagel, E., Pedersen, O.M., and Rasmussen, J.
 1981 Notes on Human Performance Analysis. RISO-M-2285.
 RISO National Laboratory, Roskilde, Denmark.
- Keeney, R.L., and Raiffa, H.
 1976 Decisions With Multiple Objectives. New York: John Wiley and Sons.

Moray, N.

- 1981 The role of attention in the detection of errors and the diagnosis of failures in man-machine systems. In J. Rasmussen and W.B. Rouse, eds., Human Detection and Diagnosis of System Failure. New York: Plenum Press.

Moray, N., Neville P., and Lootsteen, P.

- 1983 Acquisition of process control skills. Proceedings of 16th Annual Meeting of Human Factors Association of Canada: 11-14.

Moray, N., Sandersen, P., Shiff, B., Jackson, R., Kennedy, S., Ting, L.

- 1982 A model and experiment for the allocation of man and computer in supervisory control. Proceedings of the IEEE International Conference on Cybernetics & Society: 354-358.

National Research Council

- 1983 Research Needs for Human Factors. Committee on Human Factors. Washington, D.C.: National Academy Press.

North Atlantic Treaty Organization

- 1983 Conference on Human Error. Proceedings. Bellagio, Italy.

Pew, R.W., Miller, D.C., and Seehrer, C.E

- 1981 Assessment of Factors Influencing Operator Decisions. EPRI NP-1982.

Raiffa, H.

- 1968 Decision Analysis. Reading, Mass: Addison-Wesley.

Rasmussen, J.

- 1983 Skills, rules, and knowledge; signals, signs, and symbols; and other distinctions in human performance models. IEEE Transactions. SMS. SMC-13, (3) May-June.

Rasmussen, J., and Lind, M.

- 1981 Coping With Complexity. RISO-M-2293. Paper presented at the First European Conference on Human Decision Making and Manual Control, Delft, Netherlands.

Rouse, W.

- 1980 Experimental studies and mathematical models of human problem solving performance in fault diagnosis tasks. Pp. 199-216 in J. Rasmussen and W. Rouse, eds., Human Detection and Diagnosis of System Failures. New York: Plenum Press.

- Sheridan, T.B.
 1984a Supervisory control of remote manipulators, vehicles, and dynamic processes. Pp. 49-137 in W.B. Rouse, ed., Advances in Man-Machine Research. New York: JAI Press.
- 1984b Interaction of Human Cognitive Models and Computer Based Models and Supervisory Control. MIT Man-Machine Systems Lab. Report, Cambridge, Mass.
- Sheridan, T.B., and Johannsen, G., eds.
 1976 Monitoring Behavioral Supervisory Control. New York: Plenum Press.
- Simon, H.A.
 1969 The Sciences of the Artificial. Cambridge, Mass: MIT Press.
- Slovic, R., and Lichtenstein, S.
 1971 Comparison of Bayesian and regression approaches to the study of information processing in judgment. Organizational Behavior and Human Performance 11:172-194.
- Stassen, H.J.G.
 1984 The EWICS man-machine guidelines. Proceedings of the Fourth European Annual Conference on Human Decision Making and Manual Control. Institute for Perception, National Defense Research Organization (TNO), Soesterberg, Netherlands.
- Toffler, A.
 1980 The Third Wave. New York: William Morrow.
- Towill, D.R.
 1974 A model for describing process operator performance. Pp. 178-185 in E. Edwards and F.P. Lees, eds., The Human Operator in Process Control. London: Taylor and Francis.
- Tulga, M.K., and Sheridan, T.B.
 1980 Dynamic decisions and workload in multi-task supervisory control. IEEE Transactions on Systems, Man and Cybernetics SMC-10 (5) May.
- Warfield, J.N.
 1973 Intent structures. IEEE Transactions on Systems, Man, and Cybernetics SMC-3: 133-140.
- Wickens, C.D., and Kessel,
 1981 Failure detection in dynamic systems. Pp. 155-170 in J. Rasmussen and W.B. Rouse, eds., Human Detection and Diagnosis of System Failure. New York: Plenum Press.

Wood, W.T.

- 1982 The Use of Machine Aids in Dynamic Multi-Task Environments: A Comparison of an Optimal Model to Human Behavior. SM thesis, Massachusetts Institute of Technology.

Woods, D.D.

- 1982 The Cognitive Coupling of Man and Machine. Paper presented at Second International Workshop on Safety and Reliability of Industrial Computer Systems, West Lafayette, Indiana.

- 1984 Some results on operator performance in emergency events. Ergonomic Problems in Process Operations. New York: Pergamon Press.

Woods, D.D., Wise, J.A., and Hanes, L.F.

- 1982 Evaluation of Safety Parameter Display Concepts. NP-2239. Palo Alto, Calif.: Electric Power Research Institute.

Zadeh, L.A.

- 1965 Fuzzy sets. Information and Control 8:338-353.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Research and Modeling of Supervisory Control Behavior: Report of a Workshop		5. TYPE OF REPORT & PERIOD COVERED Technical Report
7. AUTHOR(s) Thomas B. Sheridan and Robert T. Hennessy, eds.		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Committee on Human Factors National Research Council Washington, D. C. 20418		8. CONTRACT OR GRANT NUMBER(s) N00014-81-C-0017
11. CONTROLLING OFFICE NAME AND ADDRESS Engineering Psychology Programs Office of Naval Research Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Work Unit No. NR196-167
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE November 1984
		13. NUMBER OF PAGES
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release: Distribution Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES The Committee on Human Factors is sponsored by the Office of Naval Research, the Air Force Office of Scientific Research, the Army Research Institute, the National Aeronautics and Space Administration, and the National Science Foundation.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Supervisory Control Model System Design Automation Man-Machine Systems Controls Mental Models Human Factors Human Performance Committee on Human Factors Simulation		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The role of the human operator is expanding to include supervisory control of otherwise automated systems. Methods for performing research on supervisory control of these modern complex systems are identified. Techniques for enhanced communication between researchers and designers are reviewed. Recommendations for future study have been prepared.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S N 0102-LF-014-6601Unclassified
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)